UNIVERSITY of CALIFORNIA

SANTA CRUZ

THE LOGIC OF CONVENTIONAL IMPLICATURES

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

LINGUISTICS

by

Christopher Potts

June 2003

The Dissertation of Christopher Potts is approved:

Professor Geoffrey K. Pullum, Chair

Professor William A. Ladusaw

Professor James McCloskey

Frank Talamantes
Vice Provost and Dean of Graduate Studies

Contents

Abstract

The Logic of Conventional Implicatures

by

Christopher Potts

The history of conventional implicatures is rocky, their current status uncertain. I return to Grice's (1975) original definition with an eye open for novel support. I argue that, even without textbook examples such as *therefore* and *but*, conventional implicatures would still be widely attested in natural language. Grice's definition characterizes a class of speaker-oriented commitments that trace back to individual lexical items and invariably yield semantic multidimensionality. These properties unify the (syntactically diverse) factual domain, which divides fairly easily into two broad classes: (*i*) supplements, including appositive relatives, nominal appositives, *As*-parentheticals, speaker- and topic-oriented adverbs, and utterance modifiers (chapter 3); and (*ii*) expressives, including adjectives like *damn*, the descriptive content of epithets, some kinds of subjunctive voice, and honorification in Japanese (chapter 4). I define a higher-order lambda calculus that provides the tools we need for formalizing Grice's definition and in turn for modelling the meanings of the expressions in (*i*)–(*ii*). The logic, which extends and sharpens the insights of Karttunen and Peters (1979), imbues the label 'conventional implicature' with theoretical content. Though considerable attention is paid to the model-theoretic aspects of the investigation, particularly as they relate to the formal modelling of discourses, much of the dissertation concerns the nature of natural language semantic composition, which we can study independently of a specific class of

structures. In the setting of the logic I define, conventional-implicature content is often distinguished solely in the meaning language. Thus, the facts under discussion seem to provide reason to view a representational language for meanings as an essential part of semantic theory. I close by asking what happens when we make slight revisions to Grice's definition. Removing speaker-orientation results in another rich class of semantically multidimensional constructions, including many that were originally classified as conventional-implicature contributors. I show that the meaning language defined here yields a theory of them as well.

Acknowledgements

When we discussed Grice's 'Logic and conversation' in my UCSC Semantics II class in the spring of 2002, I did my best to dodge questions about conventional implicatures. I felt unsure of how to identify them in the wild, so I hoped that a bit of talk about *but* would suffice. But the students (particularly Laura Buhl) persisted, in that mischievous way that only intelligent undergraduates have. I went off to try to make sense of the scene myself. A year later, I feel prepared to give the lecture they were asking for.

I did not reach this point alone. It was actually Bill Ladusaw who first pushed me in the direction of conventional implicatures, in that mischievous way that only Bill has. His insights and advice led me to some of the major proposals of this dissertation, and I'll always be grateful for his guidance on how best to interpret and present the material.

Throughout my time as a graduate student, I benefitted from the wide-ranging expertise of Jim McCloskey. During discussions with him, I enjoyed some unforgettable 'A ha!' moments, moments that marked major turning points in my work. For these (among many others), I owe him my deepest thanks.

Ash Asudeh read and commented upon a draft of this dissertation with a mix of encouragement and incisive criticism — the kind of feedback that is most useful to an author. His help led directly to major improvements. Among the many reasons I am glad this work is complete is that I can now turn my attention to the collaborative work he and I have been planning for some time.

I feel tremendously lucky to count Line Mikkelsen as a friend and colleague (800+ e-mails and counting). I will miss having her as an officemate and Pogonip trailmate. It's amazing how much intellectual ground we've covered while walking.

This dissertation was greatly improved by discussions with David Adger, Luis Alonso-Ovalle, Kent Bach, Chris Barker, Ben Caplan, Sandy Chung, Patrick Davidson, Donka Farkas, Kai von Fintel, Lyn Frazier, James Isaacs, Pauline Jacobson, Shigeto Kawahara, Chris Kennedy, John Kingston, Angelika Kratzer, Afton Lewis, Sally McConnell-Ginet, Jason Merchant, Øystein Nilsen, Stan Peters, Gillian Ramchand, Tom Roeper, Uli Sauerland, Barry Schein, Ken Shan, Yael Sharvit, Peggy Speas, Roly Sussex, Ida Toivonen, and audiences at UConn, *NELS 33*, UMass, Amherst, and USC.

I became a linguist as an NYU undergraduate, studying under Paul Postal. I hope that my work continues to show Paul's influence. Larry Hutchinson also helped to set me off in the right direction (as did the gift of a library of linguistics books from Larry and Deirdre Wheeler). Daniel Büring and Jorge Hankamer were early and important influences for me in Santa Cruz. Sandy Chung can always be counted on for tough questions — and continued encouragement when no answers are forthcoming.

It is certain that I would not be without the help of my parents. I am grateful to them, to my sister Ali, and to my Grandma Nancy for their support and encouragement. They are all teachers in one way or another. They know what's important.

Finally, I extend my most heartfelt thanks to my advisor, Geoff Pullum. I began angling for a chance to work with Geoff the moment I arrived in Santa Cruz. I soon learned that big, ambitious projects get his attention, so I bothered him continually with major research proposals until he took me under his wing. Geoff introduced me to many of the concepts and ideas that are now at the heart of my intellectual life. No one who did not run the Big Sur Marathon in 2003 has had a more important, lasting influence on me as a thinker and human being.

Chapter 1

A preliminary case for conventional implicatures

1.1 A fresh look at an old definition

The history of conventional implicatures is rocky, their current status uncertain. So it seems wise to return to their source and start fresh, with an open-minded reading of the original definition (Grice 1975) and an eye open for novel factual support. Suppose the textbook examples (*therefore*, *even*, *but* and its synonyms) disappeared. Where would conventional implicatures be then? This dissertation's primary descriptive claim is that they would still enjoy widespread factual support. I match this with a theoretical proposal: if we move just a few years forward from the genesis of CIs, we find in Karttunen and Peters' (1979) multidimensional semantics the basis for an ideal description logic.

A successful case for CIs is sure to be hard won, since their very existence has been questioned. Bach (1999) mounts a direct assault on the usual factual basis; many other working semanticists use the term 'conventional implicature' synonymously with 'presupposition', implicitly denying the need for a descriptive or theoretical distinction. Thus, it is crucial to present a rich body of evidence for CIs. But it is equally important to accompany this evidence with arguments that no other classification suffices. These are the descriptive tasks at hand, and they rightly occupy much of the present work.

The description is informed throughout by a type-driven multidimensional semantic translation language, the basis for my theory of CIs. I venture that part of the reason CIs have received short shrift from semanticists is that few attempts have been made to provide a theoretical implementation. A linguist studying presuppositions, conversational implicatures, or intonational meanings has a wealth of description logics available to use and assess. But 'conventional implicature' is usually just a label. It lacks bite; only when supported by a logical system can it be said to make predictions. For this reason, chapter 2 is pivotal. It provides an explicit description logic and relates it to familiar concepts from current semantic theories.

The evidence for CIs is drawn from diverse areas of natural language semantics, roughly divisible into two superclasses: supplemental expressions (appositives, parentheticals) and expressives (e.g., epithets, honorifics). I provide some representative examples in (1.1).

(1.1) **supplements**

a. "I spent part of every summer until I was ten with my grandmother, who lived in a working-class suburb of Boston."[1]

(*supplementary relative*)

b. "After first agreeing to lend me a modem to test, Motorola changed its mind and said that, amazingly, it had none to spare."[2]

(*speaker-oriented adverb*)

(1.2) **expressives**

a. "We bought a new electric clothes dryer, and I thought all there was to it was plugging it in and connecting the vent hose. Nowhere did it say that the damn thing didn't come with an electric plug!"[3]

(*expressive attributive adjective*)

b. saami ha-l-maʒduub nəse l-mawʕad
Sami 3-*the-idiot*.SM *forgot*.3SM *the-appointment*
'Sami, this idiot, forgot the appointment.' (Aoun et al. 2001:385, (37a))

(*Lebanese Arabic epithet*)

c. Yamada sensei-ga o-warai-ni nat-ta.
Yamada teacher-NOM HON-*laugh*-DAT *be*-PERF
'Professor Yamada laughed.' (Shibatani 1978:54, cited in Toribio 1990:539)

(*Japanese subject honorific*)

Each construction provides a novel setting in which to explore the CI hypothesis and evaluate syntactic and semantic alternatives. Two general notions unify the factual domain: these meanings are *speaker-oriented entailments* and *independent of the at-issue entailments*. I use 'at-issue entailment' as a coverterm for regular asserted content ('what is said', in Grice's terms). 'At-issue entailment' sets up a useful contrast with CIs, which are secondary entailments that cooperative speakers rarely use to express

[1] Thurman, Judith. Doing it in the road. *The New Yorker*, June 10, 2002 (p. 86).
[2] <http://www.hamline.edu/apakabar/basisdata/1997/03/21/0066.html>
[3] <http://jjdavis.net/blog/arc20010325.html>

controversial propositions or carry the main themes of a discourse. Rather, CI expressions are used to guide the discourse in a particular direction or to help the hearer to better understand why the at-issue content is important at that stage. For instance, if I utter (1.3), my primary intention is to arrive at an information state that entails the truth of the proposition that Ed's claim is highly controversial.

(1.3) Ed's claim, which is based on extensive research, is highly controversial.

With the CI content expressed by the supplementary relative, I provide a clue as to how the information should be received. This example is felicitous in a situation in which, for example, I want to convey to my audience that the controversy should not necessary scare us away from Ed's proposal — after all, it is extensively researched. Or I might use the example with a group of detractors from Ed's claim. Then the supplementary relative could indicate that we cannot expect to dispel Ed's claim solely on the basis of its controversial nature.

Expressive content is used in essentially the same way. Though we will see that expressives and supplements differ in important respects, their discourse functions are closely related. For instance, (1.4) would take us to an information state in which the speaker must mow the lawn; *damn*'s content would let us know that the speaker is displeased by this obligation.

(1.4) I have to mow the damn lawn.

Thus, we learn not only that the speaker must mow the lawn, but that we would do well not to reply with an earnest and sincere "Oh, I envy you"; this use of *damn* suggests that the discourse should head in a direction in which lawn-mowing is viewed negatively.

The effect that *damn* has on this discourse is approximately that of the supplementary relative in *I have to mow the lawn, which I hate doing.* Since the logic presented in chapter 2 assigns supplements and expressives the same kind of composition scheme, we have all the tools we need to state these pragmatic generalizations precisely.

Notably, no highlighted expression in the class represented by (1.1)–(1.2) makes a nontrivial at-issue contribution. For instance, removing *damn* from (1.2a) has no effect on the at-issue proposition expressed by its final sentence. I believe this is nonaccidental, and so formulate it as a generalization to be captured by the description logic for CIs:

(1.5) No lexical item contributes both an at-issue and a CI meaning.

This is likely to prove controversial. It excludes *but* from the class of CIs, since that item has the at-issue meaning of *and* and a purported CI dimension contrasting two properties. But Bach (1999) make a compelling case that *but* has entirely at-issue content. I strengthen the case in section 1.4.5 below; chapter 6 addresses in greater depth the kind of multidimensional meaning that *but* determines.

This introductory chapter begins with a critical overview of the early history of conventional implicatures (section 1.2). I then move to a series of brief introductions to the constructions that motivate CIs in this dissertation. Because they are little studied at present, it is necessary to establish some terminology and justify certain distinctions among expression-types. These brief remarks enrich section 1.4, which moves systematically through the kinds of meaning in current theories, supporting a range of descriptive generalizations that justify the independence of CIs from all of them.

The formalization begins in chapter 2 with the presentation of a lambda calculus with diverse enough types to isolate CIs.

1.2 A brief history

Conventional implicatures were born into neglect. Grice (1975) advanced the term and a definition, but only so that he could set such meanings aside. In ‘Logic and conversation’, he is concerned to derive conversational implicatures from the cooperative principle and the maxims of conversation. In an early passage, he acknowledges that CIs fall outside the bounds of this pragmatic theory, in large part because they represent arbitrary features of individual lexical items. There is thus no hope of calculating their presence or nature based on general principles of cooperative social interaction. The passage’s main purpose is to dispose of a class of meanings that he wishes not to discuss.

The passage is nonetheless potentially exciting. As I said, it draws a restrictive boundary around (Grice’s (1975)) pragmatic theory, placing CIs squarely outside of it. More importantly, it provides some terminology for talking about a class of expressions that permit speakers to comment upon their assertions, to do a bit of editorializing in the midst of asking questions and imposing demands. Such expressions are bound to be significant, both for what they tell us about how natural language semantic theory should look and for what they can tell us about how speakers use their languages. Thus, it is worth pulling the passage apart, literally and conceptually, in order to isolate its main insights. The following series of quotations comes from a single paragraph early in Grice 1975.

(1.6) "In some cases, the conventional meaning of the words used will determine what is implicated, besides helping to determine what is said." (p. 44)

The phrase "the conventional meaning of the words" is the crux of this statement, since it locates CIs in the lexicon. The 'conventional' part of 'conventional implicature' stands in for 'not calculable from the conversational maxims and the cooperative principle'. This is initial (and compelling) motivation for a dividing line between the phenomena that pragmatic principles should cover (conversational implicatures) and those that they cannot (CIs, among others).

One can and should refine (1.6): many expressions harbor content that does not reduce to at-issue entailments, presuppositions, intonational meanings, or conversational implicatures. The name 'conventional implicature' is a useful coverterm for this more articulated characterization, and Grice's description matches the facts discussed here. Though the descriptive and logical work of this dissertation is easily divorced from Grice's terminology, I maintain that the connection is important. His definition is essential to seeing that the constructions discussed in this dissertation are unified in significant ways. While I believe that the 'implicature' part of 'conventional implicature' is unfortunate, and that Grice failed to locate the proper factual basis for CIs, this dissertation's central themes nonetheless trace back to his work.

Grice's passage continues with an example:

(1.7) "If I say (smugly), *He is an Englishman; he is, therefore, brave*, I have certainly committed myself, by virtue of the meaning of my words, to its being the case that his being brave is a consequence of (follows from) his being an Englishman." (p. 44)

This passage again sets up a contrast with conversational implicatures. CIs are entailments, whereas conversational implicatures are not — they are context-dependent and always negotiable. If I say, “Eddie has three bicycles”, I conversationally implicate that, for any n greater than 3, it is false that Eddie has n bicycles. But this is not a commitment; I could felicitously continue, “Hey, let’s be honest: Eddie has ten bicycles. He’s a bike junky”. CIs permit no such cancellation; following any of the sentences in (1.1)–(1.2) with a denial of the content of the highlighted expression results in an incoherent discourse. In this sense, CIs pattern with at-issue entailments.

But Grice takes steps to distinguish CIs from at-issue entailments as well:

(1.8) “But while I have said that he is an Englishman and said that he is brave, I do not want to say that I have *said* (in the favored sense) that it follows from his being an Englishman that he is brave, though I have certainly indicated, and so implicated, that this is so. I do not want to say that my utterance of this sentence would be, *strictly speaking*, false should the consequence in question fail to hold.” (p. 44–45)

In no uncertain terms, Grice defines CIs as disjoint from at-issue entailments. I take seriously the intuition expressed in (1.8), which is a suitable articulation of the uneasiness one has about the semantic value of an utterance containing a false or inappropriate conventional implicature. Once it is accepted (following Jackendoff (1972), Bellert (1977), and Bach (1999)) that some sentences can express multiple, nonconjoined propositions, Grice’s intuition can be made precise and formally implemented.

The passage also relativizes CI content to the speaker of the utterance in question (this is noted also by Chierchia and McConnell-Ginet (1990:§6.4.3), though they read the passage somewhat differently). This too is a significant aspect of the argument for CIs. All genuine examples of CIs involve a contribution that the speaker makes to an utterance. They are speaker-oriented comments on a semantic core (at-issue entailments); we so effortlessly separate the two kinds of meaning that these comments can, and often do, appear in accurate indirect speech reports, as in examples (1.1b), (1.2a), and many to come.

For better or worse, Grice (1975) drops CIs at this point:

(1.9) "So *some* implicatures are conventional, unlike the one with which I introduced this discussion of implicature.

I wish to represent a certain subclass of nonconventional implicatures, which I shall call *conversational* implicatures [. . .]." (p. 45)

Grice moves to a discussion of conversational implicatures, one that continues to form the backbone of work in pragmatics. He planted the seed for CIs and then moved on. Since then, numerous proposals for CIs have been offered. But few if any have shown much staying power, though for reasons that do not impact the importance of this class of meanings, but rather only the appropriateness of the evidence brought to bear on the subject.

As I said, I aim to do without textbook examples like *therefore.* So, in (1.10), I extract the abstract properties of CIs from the above series of quotations.

(1.10)
a. CIs are part of the conventional (lexical) meaning of words.
b. CIs are commitments, and thus give rise to entailments.
c. These commitments are made by *the speaker of the utterance* "by virtue of the meaning of" the words he chooses.
d. CIs are logically and compositionally independent of what is "*said* (in the favored sense)", i.e., independent of the at-issue entailments.

I return to this definition often; it is the job of chapter 2 to show how these clauses translate into a multidimensional description logic, with most of the work done by the presence of independent dimensions of meaning.

Chapter 2 is rooted in the ideas of Karttunen and Peters (1979). It is worth stressing, then, that the definition of 'conventional implicature' adopted by them differs radically from (1.10). It has often been noted that Karttunen and Peters actually give a logic for presuppositions, not conventional implicatures as Grice understood them. That they perform this terminological slight of hand is evident from their descriptive characterization of 'conventional implicature':

(1.11) "As a general rule, in cooperative conversation a sentence ought to be uttered only if it does not conventionally implicate anything that is subject to controversy at that point in the conversation. Since the least controversial propositions of all are those in the common ground, which all participants already accept, ideally every conventional implicature ought to belong to the common set of presumptions [. . .]." (p. 14)

This is an injunction that CIs be backgrounded, i.e., entailed by the shared knowledge of the discourse participants at the time of utterance. But nothing in the above series of quotations from Grice 1975, on even the most creative of readings, suggests that CIs should be backgrounded. As van der Sandt (1988) writes, "Karttunen and Peters do not make it clear why conventional implicatures belong to the common ground. This view certainly cannot be found in Grice and as far as I know has not been argued for elsewhere in the literature." (p. 74). The backgrounding requirement is one of the central descriptive properties of presuppositions; it is so wrapped up with this notion that Karttunen and Peters (1979) even use 'presumption', an alternative to 'presupposition', in the above quote.

Their confounding of 'conventional implicature' and 'presupposition' happens at the theoretical level as well. Karttunen and Peters' (1979) grammar fragment employs a 'heritage function' to regulate how CIs interact with higher operators. It becomes clear upon inspection that this is a new name for a presupposition projection function. I discuss this issue again in chapter 2, section 2.10, under the rubric of the 'binding problem'. Suffice it to say here that I require no heritage function. Including one in the theory would obscure the important fact that CIs always project to the highest possible point.

Karttunen and Peters intended their redefinition of 'conventional implicature' to clarify the theory of non-at-issue content. But it had the opposite effect: we lost sight of Grice's definition. For many authors, 'conventional implicature' and 'presupposition' are interchangeable, despite the fact that (1.10) barely resembles the usual definition of 'presupposition'. In turn, Karttunen and Peters 1979 is generally regarded as a theory of presuppositions. Cooper (1983), Heim (1983), Beaver (1997, 2001), Krahmer (1998), and Dekker (2002) all adopt this interpretation of the work. The following passage from Gamut 1991, an introductory textbook in logic and linguistics, is typical of the way this issue is negotiated.

(1.12) "Karttunen and Peters [...] proposed translating natural language sentences ϕ as pairs of formulas $\langle \phi^t, \phi^p \rangle$, in which ϕ^t represents ϕ's truth conditions and ϕ^p represents its presuppositions (or what they call *conventional implicatures*)." (p. 188)

Later, the authors are more deferential to Karttunen and Peters' terminology, but this quotation gets right at the heart of the quiet shift that Karttunen and Peters performed. I emphasize that I do not follow them in the redefinition. It is more fruitful to explore (1.10), which identifies a new class of meanings. The pressing question is whether these meanings are attested in natural language.

1.3 Factual support for CIs

This section introduces the constructions that play a leading role in chapters to come. I review some basics of their syntax and describe their interpretive properties in a general way, in preparation for detailed study later. This is also a chance to show briefly how each contributes a vital element to the overall picture.

Roughly speaking, the constructions divide into two groups: supplemental (appositive) expressions, including supplemental clauses and supplemental adverbs, and expressives. From a syntactic perspective, this is a mixed crew. But the constructions are united in contributing discourse-new, speaker-oriented entailments — CIs.

1.3.1 Supplemental expressions

Supplements (appositives, parentheticals) are the finest advertisement for the CI hypothesis known to me. Though Grice seems not to have had them in mind when defining CIs, the clauses of (1.10) pick out the highlighted constructions in (1.13) unambiguously.

(1.13) a. Ames was<u>, as the press reported,</u> a successful spy. (*As-clause*)

b. Ames<u>, who stole from the FBI,</u> is now behind bars. (*supplementary relative*)

c. Ames<u>, the former spy,</u> is now behind bars. (*nominal appositive*)

Many of the important properties of these expressions turn up also in *not even* tags (*Ed didn't show up, not even for the end*), niched conjunctions (*Luke has — and you'll* *never believe this — eaten fifty eggs*), and a host of other clausal appositives. I largely restrict my attention to the constructions in (1.13), favoring depth of coverage over breadth. It is a mistake to treat all supplements, even all clausal ones, as though they were the same construction. Chapter 3 identifies numerous nontrivial ways in which the constructions in (1.13) differ from one another.

In addition to the clausal supplements in (1.13), I study a host of parenthetical adverbs, including the *speaker-oriented* and *topic-oriented* adverbs exemplified in (1.14).

(1.14) a. {Cleverly/Wisely}, Beck started his descent. (*topic-oriented adverbs*)

b. {Unfortunately/Luckily}, Beck survived the descent. (*speaker-oriented adverbs*)

What I call 'topic adverbials' are the 'subject-oriented' adverbs of Jackendoff (1972) and Bellert (1977). I eschew the old term because it wrongly suggests that these items invariably predicate something of the grammatical subject. In truth, the entity-level argument is often merely a salient discourse topic, as in (1.15), in which the agent characterized as thoughtful seems to be the keyboard's designers.

(1.15) "Physically, the keyboard is smaller than I expected, and extremely well built — there's no creaking or flexing. The keys look as if they will last well — including their paint. Thoughtfully, there is a clip-on cover for the connector while not in use."[4]

Supplements have much to offer the theory of CIs. Unlike some of the other expressions discussed in this dissertation, it is straightforward to determine their propositional contribution, which is given in the expected way by the internal structure of the supplement and its main clause adjunction point. For example, in virtue of being adjacent

[4] <http://www.pdatweaks.com/reviews.php?itemid=238>

to (i.e., the sister of) *Ames* in (1.13b), the supplementary relative *who stole from the FBI* contributes the proposition that Ames stole from the FBI. Once entered into the context, this proposition behaves like any other: it can be pronominalized with *do so* and similar elements; it can serve to license additive modifiers like *also*; and so forth.

Facts such as these provide straightforward evidence for the claim of chapter 2 that at-issue and CI expressions can have the same models; *spy* denotes the characteristic function of the set of spies whether it is inside a nominal appositive (as in (1.13c)) or in a main clause predication. The distinction between at-issue and CI content is often entirely about semantic composition. It is properly located in the meaning language, as a syntactic fact about the logic of the natural language semantics. Thus, at the heart of this dissertation is a nontrivial appeal to a semantic translation language. While the claim sounds controversial (controversially anti-Montagovian), I show in chapter 2 that no theory of natural language syntax and semantics has managed to do without a meaning language of some kind.

There is an even more syntactic option of course: we could, following work by McCawley (1982, 1987, 1989, 1998) and Huddleston and Pullum (2002) (and many others), assign to supplements, and perhaps all CI contributors, a distinguished syntax. I explore this hypothesis in chapter 5, using it to develop a version of the *wide-scope co-ordination hypothesis* for supplements that does not run afoul of the known arguments against it. But this purely syntactic proposal leads to unacceptable redundancies in the syntactic description and yields no new benefits or insights in the semantics. It was arguably doomed from the outset, since it attempts to build a fundamentally semantic concept into the syntactic structures. The syntactic approach does, though, bring to the fore a general feature of the constructions addressed here: they seem always to lead us to multidimensional concepts. As I discuss in chapter 4, semantic non-CI analyses must also propose multidimensional sentence meanings. It is vital that one keep this shared feature in mind when assessing alternative proposals.

In my study of adverbs in chapter 3, I distinguish the adverbials exemplified in (1.14) from the *utterance modifiers* in (1.16), discussed briefly by Jackendoff (1972) and more systematically by Bellert (1977) and Bach (1999:§5).

(1.16) a. <u>Confidentially (speaking),</u> Sal is about to get canned.
b. <u>(Speaking) Just between friends,</u> Sal is about to get canned.
c. <u>Frankly (speaking),</u> Ed fled
(*utterance modifiers*)

'Utterance modifier' sounds like a semantic–pragmatic designation, but it is frequently cashed out in syntactic terms, as a functional projection (Cinque 1999). I make good on the intuition reflected in the labels 'utterance modifier', 'pragmatic adverb' (Bellert 1977), and 'second-order modifier' (Bach 1999), by analyzing these expressions in terms of what I call *discourse structures* (defined in chapter 2, section 2.8). These are *layered* in the sense that one can view them as involving a larger (upper) structure that contains a set of smaller (lower) structures. The upper layer provides a semantics for discourses and the objects they contain. The lower layer lets us talk about individual sentences and their meanings. In this setting, we can give a precise semantics for the paraphrase of (1.16c) in (1.17).

(1.17) The speaker frankly utters the sentence *Ed fled.*

The semantics for utterance modifiers is located mainly in the upper layer of the logic and model theory. This makes intuitive sense when one sees that the upper layer is a formerly metagrammatical level now brought into the grammar. We thus have a direct translation of the notion, found in traditional grammars and present-day style books (Williams 1990), that utterance modifiers belong to a kind of metalanguage that we use for talking about discourses (speeches, texts, conversations).

The same basic treatment extends to uses of utterance modifiers with interrogative complements, where the meaning that the adverb contributes is somewhat different than it is in the presence of a declarative.

(1.18) a. Confidentially, is Al having an affair?
$\approx$ I promise to keep the answer to *Is Al having an affair?* a secret.
b. Honestly, has Ed fled?
$\approx$ Provide me with an honest answer to the question *Has Ed fled?*

The adverbs function here to request something of the hearer. The meaning change is evidently engendered by the presence of an interrogative complement rather than a declarative one. The discourse structures I define are sensitive to such distinctions; a concise description of these readings is readily available.

Many have noticed that supplements do not contribute their meanings in the usual fashion, and a variety of different, disparate methods for modelling the contribution has been identified (Keenan 1971; Boër and Lycan 1976; Emonds 1976; the above-mentioned work by McCawley). The CI hypothesis, grounded in the multidimensional approach, captures what is right about all of these past proposals, but without their unpalatable consequences. The analysis I offer permits us to interpret surface structures in which supplements are syntactically embedded like regular modifiers. The well-known fact that nothing scopes over their meanings is handled in the meaning language.

1.3.2 Expressives

The characterization of CIs as comments upon a semantic core is nowhere more fitting than with expressives. Such expressions are vital to naturally occurring discourses: searching *damn* or *friggin* on the Internet turns up tens of thousands of relevant hits; honorific marking runs through essentially all discourse in languages like Japanese and Thai; and discourse particles are a notable and defining feature of German, Danish, and many other languages. This ubiquity should guarantee expressives a place in semantic and pragmatic theories. However, to date, theoretical semanticists have contributed only a handful of works on the topic. As a result (and quite happily) we still get to confront foundational questions in this domain.

Though of limited size, the literature on expressives converges on a few essential concepts. The semantic multidimensionality of sentences containing expressives is brought to the fore and given a preliminary technical interpretation by Kratzer (1999). Expressives' speaker-orientation is noted by Cruse (1986:271ff) and Löbner (2002:§2.3). The expressive attributive adjective (EA) in (1.19) corroborates their observations.

(1.19) "We bought a new electric clothes dryer [. . .] Nowhere did it say that the <u>damn</u> thing didn't come with an electric plug!"[5]

The expressive is inside an indirect quotation, and yet its content is independent of

[5] <http://jjdavis.net/blog/arc20010325.html>

whatever meaning is the argument to the higher predicate, and in turn to any other argument. The speaker of (1.19) makes manifest his heightened emotions, and yet we intuit that neither the frustration nor the speaker's emotive contribution is included in the instructions for the clothes dryer (the meaning of *it*). These observations together exemplify clauses (1.10c) and (1.10d) (CIs' speaker-orientation and independence of the at-issue content, respectively). In virtue of tracing back to *damn*, the expressive content satisfies the lexicality property (1.10a). Finally, the commitment property (1.10b) is clearly on display: some expressives are so powerful that speakers cannot even use them in jest without committing themselves to their content.

As reviewed in section 1.4, the invariance of this content under the presupposition plug *say*, and the related but distinct fact that it must be teased apart from the proposition expressed by the main clause, both indicate that this content is neither a presupposition nor an at-issue entailment. The fact that we can locate the relevant meaning in a specific lexical item tells against a treatment in terms of conversational implicatures. The content's invariance under negation, tense, modalization, questioning, and conditionalization, as well as its general noncancellability, speak decisively against this classification. In sum, EAs are prime candidates for a CI analysis.

EAs are perhaps best thought of as a special class of attributive adjectives that can never contribute at-issue content. Many adjectives seem to alternate between at-issue and CI readings. Adjectives with objective truth-conditions (*red*, *Swedish*) are likely to hide this dimension, but it is evident with, for example, *lovely* in (1.20).

(1.20) Edna is at her friend Chuck’s house. Chuck tells her that he thinks all his red vases are ugly. He approves of only the blue ones. He tells Edna that she can take one of his red vases. Edna thinks the red vases are lovely, selects one, and returns home to tell her housemate,

“Chuck said I could have one of his <u>lovely</u> vases!”

Edna characterizes Chuck’s vases as lovely. The adjective is nonrestrictive, and it is not part of what Chuck said to Edna. If it were, then *his lovely vases* would denote the set of Chuck’s *blue* vases. But Edna was not licensed to take any of them. We easily recognize that Edna is contributing the adjective; the utterance expresses two propositions: (*i*) that Chuck said Edna could have one of his vases; and (*ii*) Edna thinks Chuck’s vases are lovely. The second of these is CI content.

The special value of these cases is that they display a minimal deviation from the expected isomorphism between the syntax and the semantics: the meaning of *lovely* does not take the meaning of *vases* as its argument. Rather, the composition scheme involves **lovely** applying to the entity-level term **vases-of**(**chuck**), which is not the meaning of a surface syntactic constituent in (1.20).

This mismatch between the syntactic structure and the semantic composition is a controlled form of the variability that EAs display. Though nominal-internal, EAs can take common nouns, full nominals, and full clauses as their arguments. The examples in (1.21) can be interpreted in a way that brings out each of these readings.

(1.21)
- a. I have seen most bloody Monty Python sketches!
  (*the speaker disapproves of Monty Python sketches in general*)
- b. I hate your damn dog! (It’s not nearly so friendly as my dog.)
  (*the speaker disapproves of the addressee’s dog*)
- c. My friggin’ bike tire is flat again!
  (*the speaker disapproves of the fact that his bike tire is flat again*)

Syntactic movement of English attributive adjectives is contraindicated by all known syntactic tests. Hence, we must call upon the semantics to ensure that the meaning of *damn* can apply to noun-phrase and clausal meanings (at least) despite its nominal-internal position in the syntax. As a result, we obtain additional arguments that semantic representations play a nontrivial role. As noted above, the at-issue/CI divide is located in the meaning language. A direct mapping from natural language expressions to model-theoretic objects erases the distinction and hence does not suffice. The lambda terms themselves are an essential stopping off point. The interpretive properties of these attributive adjectives provide additional evidence that we interpret something more articulated than mere surface strings.

If one thinks in semantic terms, one expects to find common nouns with expressive (CI) meanings, given the semantic similarities between adjectives and common nouns. The expectation is met; theoretical linguists call such nominals *epithets*. Informally speaking, epithets are pronouns with some added punch, in the form of emotive descriptive content. They are often called upon in the syntactic literature as evidence for or against particular views of the syntactic binding theory. But their semantics is relatively unexplored. A notable exception is Asudeh (2003), where epithets are discussed from the perspective of Glue semantics. Chapter 2, section 2.12, is a discussion of the points of contact between Asudeh's theoretical framework and those of the present dissertation.

Like EAs, epithets can appear inside indirect quotations without forming part of the semantic content of the reported utterance. I illustrate with the instance of donkey anaphora in (1.22), which I owe to Ash Asudeh (p.c., 6/02). (Here and throughout, I use numerical subscripts to indicate semantic binding relationships.)

(1.22) Every Democrat advocating [a proposal for reform]$_1$ says [the stupid thing]$_1$ is worthwhile.

We can use epithets to illustrate each of the clauses in the definition of CIs in (1.10). It is clear that we should locate the expressive content on the epithet; this satisfies the lexicality condition (1.10a). The sentence involves the speaker's characterization of Democratic proposals for reform as stupid; the truth of this sentence does not require that every Democrat characterize his proposal as both stupid and worthwhile. Nor need every Democrat recognize that the speaker views these proposals as stupid. Thus, the epithet's contribution is independent of the at-issue proposition (expressible by substituting a pronoun for *the stupid thing* in (1.22)). So the grammar must separate these meanings — they are intuitively independent (clause (1.10a)).

It is imperative that this separation happen only at the level of meanings; epithets are syntactically integrated (often as argument nominals). The CI approach developed in chapter 4 achieves this result. Its only published competitor to date is Kaplan's (1989:555, fn. 71) brief suggestion that quantifying-in is an appropriate mechanism. Chapter 4, section 4.5, shows that quantifying-in alone does not yield an accurate description even when supplemented with numerous ad hoc premises.

So epithets are another area of support for CIs. But they offer much more. More than any other construction, they test the limits of the description logic and its treatment of quantification. As noted, *the stupid thing* in (1.22) is a donkey pronoun; on Heim's (1982) classic analysis of donkey sentences, *every Democrat advocating a proposal for reform* quantifies over Democrat–proposal pairs. At first, this seems naturally represented as in (1.23).

(1.23) Every Democrat advocating [a proposal for reform]$_1$ says [the stupid thing]$_1$ is worthwhile.

at-issue:

$$\forall\langle x,y\rangle\left(\left(\begin{array}{rl}\mathbf{democrat}(x) & \wedge\\ \mathbf{ref\text{-}proposal}(y) & \wedge\\ \mathbf{advocate}(y)(x) & \end{array}\right)\rightarrow\mathbf{say}(\mathbf{worthwhile}(y))(x)\right)$$

CI: $\mathbf{stupid}(y)$

But in the description logic of chapter 2, as in Karttunen and Peters 1979, the occurrence of y in the CI dimension is not bound by the diadic universal in the at-issue meaning. One way to phrase this is that the logic inherits the 'binding problem' that Karttunen and Peters (1979:53) recognize in their two-dimensional logic. But calling it a problem is misleading. It is a feature of an internally consistent logic. It could only become a (linguistic) problem if it failed to describe some natural language facts that it was intended to describe. The reverse seems to be the case; chapter 2, section 2.10, reviews evidence from supplemental expressions that this limitation can be a virtue. But in extending the analysis to epithets, do we lose those results?

The answer is that we clearly do not. On the contrary, inspection of a broader range of cases — in particular those that do not involve universal quantifiers — reveals that, in (1.22), the relationship between the quantifier *every* and the expressive content of *stupid thing* is not one of binding. Rather, what we seek for the expressive meaning in quantified cases is a generic quantification over the restriction on the at-issue quantifier. For (1.22), our target CI meaning is roughly 'in general, Democratic proposals for reform are stupid'. If we adopted (1.23) as representative of the translation procedure and in turn adjusted the logic so that $\forall\langle x,y\rangle$ directly linked with the CI dimension, then we would end up with an analysis that badly mishandled the data.

EAs and epithets are similar semantically, so it seems wise to show that the basic techniques developed for them extend to other, more diverse items that seem classifiable as expressives. Chapter 4 therefore closes with CI-based analyses of honorifics in Japanese and the discourse subjunctive in German, henceforth *Konjunktiv I*.

The Japanese honorific system is extremely intricate. I do not attempt complete descriptive coverage. My strategy is to concentrate on two subtypes of honorific marking: verbal honorific marking indicating the speaker's relation to the grammatical subject, as in (1.24), and performative honorific marking ('polite speech'), as in (1.25).

(1.24) Yamada sensei-ga o-warai-ni nat-ta.
Yamada teacher-NOM HON-*laugh*-DAT *be*-PERF
'Professor Yamada laughed.'
(Shibatani 1978:54, cited in Toribio 1990:539)
(*subject honorific*)

(1.25) Ame ga huri-masa-ta.
rain SUBJ *fall*-HON-PAST
'It rained.' (Harada 1976:502)
(*performative honorific*)

For subject honorification, we need to establish a connection between a morpheme on a matrix verb and that verb's subject argument. For performative honorification, we face the sort of puzzling unembeddability that is a hallmark of utterance-modifying parenthetical adverbs like *frankly* and *confidentially*, discussed briefly above and in detail in chapter 3.

The German Konjunktiv I is useful in heading off a presuppositional alternative. Briefly, Konjunktiv I is used on the inflected verb in a clause C to indicate that the speaker is not committed to the truth of the proposition expressed by C. The examples in (1.26), in which KONJ indicates Konjunktiv I morphology, help clarify this meaning contribution.

(1.26) a. Sheila behauptet, dass sie krank sei.
*Sheila maintains that she sick be.*KONJ
'Sheila maintains that she is sick.'

b. #Ich behaupte, dass Sheila krank sei.
*I maintain that Sheila sick be.*KONJ
'I maintain that Sheila is sick.'

The second example is marked because it imposes contradictory demands. The at-issue assertion is that the speaker maintains that Sheila is sick. The CI proposition induced by the Konjunktiv I morphology is that the speaker is not committed to the proposition that Sheila is sick. Were the Konjunktiv I morpheme a presupposition trigger, one would expect cancellation at the hands of the more rigid at-issue assertion. But this is not what we find, paving the way for a treatment using the CI logic of chapter 2, which treats at-issue and CI propositions identically with regard to the strength of speaker-commitment.

In addition to this presuppositional alternative for Konjunktiv I, chapter 4, section 4.7, explores in detail a scope-based alternative. Broadly speaking, this account is an effort to assimilate expressive content to the at-issue dimension. The basic mechanism for doing this is a stipulation that expressives have an intensional argument that must be filled by the actual world index — the argument that ultimately takes the proposition expressed by the sentence to a truth value. I show that such an approach can describe the basic facts — basically a given, since it is an extremely powerful and general idea. However, its freedom is its downfall. In assimilating EAs to modal and temporal modifiers like *former* and *potential*, it wrongly produces a wide range of ungrammatical readings, ones that the CI logic blocks without extra statements. What is more, the account must adopt a multidimensional perspective on sentence denotations, and it must include some method for marking certain items as expressives. These two moves are defining features of the CI logic $\mathcal{L}_{\mathrm{CI}}$ of chapter 2. Thus, it seems that this might not be an alternative in a substantive sense even if it were brought to a point where it properly described the facts in a rigorous way.

Thus, we arrive at a theme of this dissertation: the question is not whether a multidimensional theory is motivated — it seems inevitable — but rather how best to formalize the notion. This is a theoretical insight that I expect to survive even drastic revisions to the description logic I offer.

1.4 Kinds of meaning

There is a definite trend in formal semantics towards treating all semantic content that is not, intuitively, part of the semantics proper as though it were presupposed. 'Presupposition' is in danger of becoming a label for whatever part of the semantics eludes the analysis being offered — a new version of the old pragmatic wastebasket. However, the analogy is only partially accurate. Whereas few peered into the pragmatic wastebasket, the presupposition wastebasket is tended by ambitious theorists who are happy to accept whatever is tossed their way. The danger, of course, is that the result will stretch the notion of presuppositions too thin, erasing hope of a unified formal reconstruction of the concept. It is both strange and worrisome to encounter talk of, in effect, nonpresupposed presuppositions. This seems to obscure genuine diversity. At the very

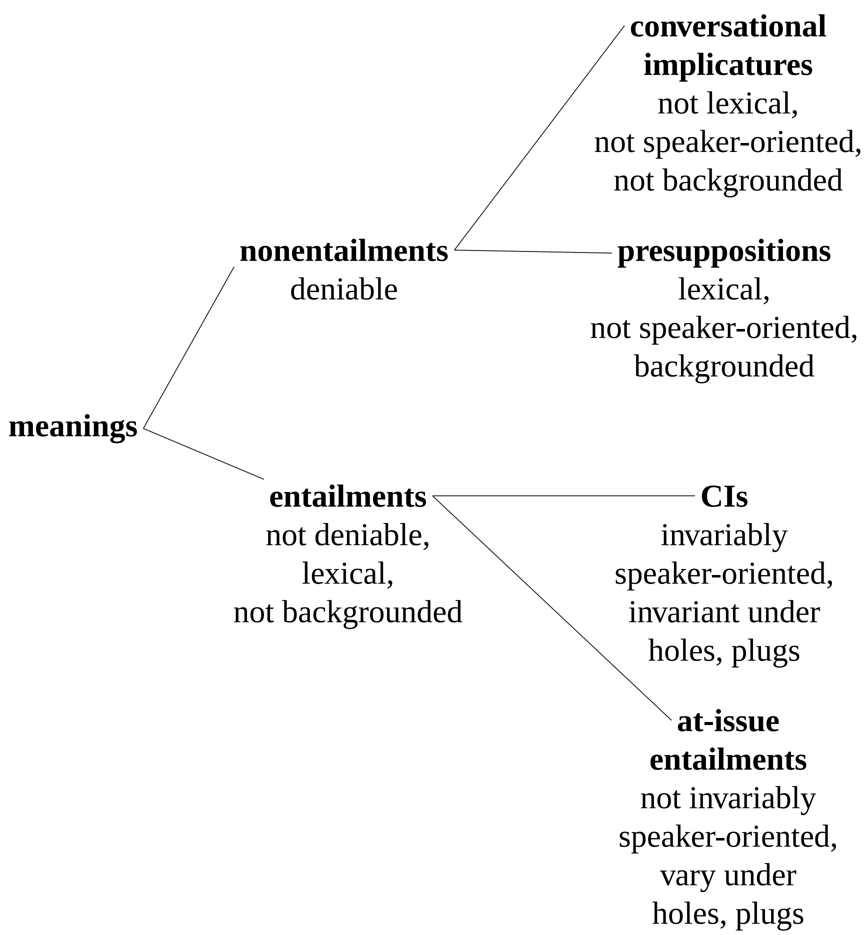

Figure 1.1: A MEANING TREE

least, we will need many subclasses of presupposition, only one of which reconstructs the pretheoretical notion of 'content that needs to be true at the point of an expression's use for that expression to mean anything'.

Thus, the backdrop for this work is a rich ontology of classes of meanings, represented in the meaning tree in figure 1.1 (page 26), which is decorated with concepts that are the focus of the next few subsections. The most inclusive class is that of *meanings* or *implications*. The meanings divide into two subclasses, *entailments* ('commitments') and *nonentailments*; the main factor in the split is the notion of *deniability*. The question, 'Is p deniable in C?' should be read as a shorthand for the question, 'Is it possible that p is a potential, but not an actual, contribution to C?' Nonentailments are deniable: it is often the case that the context conspires to alter or eliminate a potential nonentailment. In contrast, entailments are not deniable; there is no substantive distinction in this area between potential and actual meaning.

Each major subclass divides again. The nonentailments differ significantly in their origin (lexical items or general features of cooperative social interaction) and also in their discourse properties (the question of whether the meaning must be a part of the common ground at the point in the discourse at which it is offered).

At-issue entailments usually go by other names. A common term in philosophically-oriented work on the subject is 'what is said'. But this is confusing in a purely linguistic context, in which 'what is said' is likely to be equated with the complement to a verb like *say*, and might even be used to refer to the words in an utterance (rather than its semantic content), or even the pronunciation of those words. In this work, I use 'what is said' and similar phrases only in describing utterances.

In linguistics, the most common term is 'assertion' (Stalnaker 1979). But this too is not quite right; few would deny that, in (1.27), the first sentence of a published book review, the writer intends to assert that her grandmother lived in a working-class suburb of Boston. But it is wrong to treat this on par with the proposition that she spent part of every summer until she was ten with her grandmother.

(1.27) "I spent part of every summer until I was ten with my grandmother, who lived in a working-class suburb of Boston."[6]

The terminology employed here helps us to recognize that we have two assertions in (1.27). But the supplementary relative *who lived in a working-class suburb of Boston* plays a secondary role relative to the information conveyed by the main clause. The issue is not where the grandmother lived, but rather the fact that the speaker summered with her as a child. The supplementary relative's content just provides us with some important (nonlogical) consequences of this proposition — in this case, probably sociological ones inferrable from the environment she specifies.

Karttunen and Peters (1979) use the term 'extensional' for my 'at-issue entailments'. But 'extensional' is better reserved for the mode of semantics in which the interpreted structure is a first-order model, with no intensional types and propositional expressions interpreted by a set of truth values. 'Extensional' should remain a counterpoint to 'intensional', not 'implicative'.

To be sure, even 'at-issue entailment' has drawbacks. For instance, people sometimes use main clauses to say things that are not at issue in the sense that they are unresolved in the discourse. Horn (1991) seeks to make sense of these cases, and Barker and Taranto (2003) look at the specific case of the adjective *clear*. If proposals of this sort were eventually to reveal the inadequacy of using 'at-issue' in the way that I do, this would not affect the substance of the present proposal. The logic of chapter 2 does the important work, and it is of course free from this terminological morass.

The tree in figure 1.1 is partial, at least in a taxonomic sense; one can make further

[6] Thurman, Judith. Doing it in the road. *The New Yorker*, June 10, 2002 (p. 86).

distinctions. Conversational implicatures have particularized and generalized variants. Particularized conversational implicatures are highly context-dependent, whereas generalized implicatures are essentially part of speakers' knowledge of how to use language, rather than social normatives more generally. Generalized implicatures include scalar implicature and also the force of, e.g., "Do you have the time?" (which is roughly "Please tell me the time if you can"). These subclasses are not of much concern here, as they are merely labels; the move from particularized to generalized is surely gradient, and scalar implicatures do not require separate pragmatic mechanisms or behave differently from others in a theoretically significant way. They are natural consequences of a scalar semantics and the usual stock of conversational maxims.

In the remainder of this section, I review the factual considerations that place CIs on a separate branch. I do not devote much attention to justifying, for example, the distinction between presuppositions and conversational implicatures. Since both kinds of content are deniable, establishing the distinction is a delicate matter. But the literature on the usefulness of isolating presuppositions from the pragmatics is vast and rich; as Chierchia (2001:8) says, "it was thought early on that presuppositions constituted a purely pragmatic phenomenon, not amenable to a grammar driven compositional treatment [...]. But eventually it turned out that such a treatment is, in fact, the one that gets us the better understanding of the phenomenon". A recent and impressively comprehensive review of the literature is Beaver 1997. Beaver (2001) and Krahmer (1998) provide flexible, easy-to-use logical theories for presuppositions.

The task of isolating CIs would be easier if presuppositions and conversational implicatures were of a piece, or if presuppositions were distributed throughout the other classes, because there would then be fewer competing classifications. I strongly believe that such a conflation is false, and proceed on the assumption that we have at least the distinctions in figure 1.1, but skeptics of presuppositions might keep in mind that assuming a multitude of non-CI classes only makes the present job more challenging.

Intonational meanings are not easily included in the structure in figure 1.1. They function as triggers of non-at-issue meaning, but impose no further restrictions on what kind of meaning this is. I briefly address the distinctness of intonational meanings and CIs below. Showing that these are different is straightforward. But the theory of intonational meanings has much to offer the present study. It is helpful as a kind of campaign point: alternative semantics for focus is a well-accepted multidimensional view of meaning, a kind of precedent for the current study. More concretely, intonational meanings play a central role in explaining the special properties of supplements (chapter 3).

1.4.1 CIs versus conversational implicatures

Despite the occurrence of "implicature" in both names, the easiest distinction to make is between CIs and conversational implicatures. As noted above, Grice (1975) seems to have defined CIs specifically to separate them from conversational implicatures, his domain of inquiry. The definition was a way to forestall objections that the maxims leave the presence of some non-at-issue content mysterious.

The differences (listed in figure 1.1) have a common source: conversational implicatures exist in virtue of the maxims and the cooperative principle, whereas CIs are idiosyncratic lexical properties. Put another way, conversational implicatures are not inherently linguistic, whereas CIs are inherently linguistic. It is worth amplifying this point. On Grice's (1975) conception, the maxims are independent of language. He writes, "one of my avowed aims is to see talking as a special case or variety of purposive, indeed rational, behavior" (p. 47), and he in turn cites some nonlinguistic examples to illustrate how the maxims work. These comments have a technical translation: the maxims are about relations among *propositions* — model-theoretic entities that languages might pick out, but that are not grounded in language. A classic instance of a generalized conversational implicature draws out the importance of this point:

(1.28) a. "Can you pass me the salt?"
b. conversational implicature: pass the salt to me if you can

This utterance is generally interpreted as conversationally implicating that the addressee pass the salt to the speaker, so that a "yes" answer unaccompanied by an act of passing the salt is infelicitous even if true. But the presence of this conversational implicature traces back to no specific feature of the utterance; all of (1.29a–d) can convey (1.28b), and other variants are easily found. This is just the well-known 'nondetachability' of conversational implicatures: low-level tinkering with the form of the utterance is unlikely to remove such implicatures.

(1.29) a. "Are you able to reach the salt?"
b. "I could sure use the salt."
c. "My dish could use a salting."
d. "Could you send the salt my way?"

The conversational implicature generalizes along another line: in a broad range of situations,"Can you pass me X?" carries the implicature for any choice of X (*wrench*, *book*, *jodhpurs*). The unifying feature of all these cases is not a linguistic matter. Rather, we arrive at it by way of the maxims; I provide an informal calculation in (1.30).

(1.30) a. Cooperative agents do not request information they already possess. Such requests do not increase the collective knowledge of the discourse participants and so always fail to qualify as informative, relevant, and sufficiently brief.

b. If the addressee is not near the salt, then the speaker already knows that the answer to the literal readings of (1.28)–(1.29) is "no".

c. If the addressee is near the salt, then the speaker already knows that the answer to literal readings of (1.28)–(1.29) is "yes".

d. Hence, (1.28)–(1.29) must not be questions at all.

e. Some reflection on our current context suggests to the speaker that (1.28)–(1.29) must be indirect ways of asking for the salt.

The first premise is a general use of the maxims of quantity, relevance, and brevity. It is not a falsifiable principle, but rather has the status of a contractual obligation (an analogy Grice (1975:48–49) toys with). The next two premises are facts about the particular discourse. They are essential to the conclusion in (1.30d); if we remove either of them, the conversational implicature disappears. For instance, it might be felicitous to ask "Can you pass the salt?" when there is no realistic possibility that the addressee can respond with a passing action. This might occur if the addressee has recently broken both arms or is living in a society in which people of his kind are rarely permitted to touch others' foodstuff. In such contexts, neither (1.30b) nor (1.30c) is a reasonable inference. The linguistic stuff (the sentence uttered) remains the same in the two situations just described. But the nonlinguistic stuff (the collective knowledge) changes, and with it the conversational implicatures change. As a result, conclusion (1.30d) is not made, and the descriptive effect is that the usual conversational implicature is cancelled, where ‘cancelled’ is a shorthand way of saying that the discourse failed to support it in the first place.

In contrast, CIs cannot be teased apart from the lexical items that produce them, nor can we understand where they arise by appeal to the nature of the context and the maxims. The result is that they are not contextually variable. I do not here adduce evidence for this claim for each of the expressions under discussion, but it is worth illustrating with an example involving expressive modification, where cancellability might seem to be a real possibility. Suppose that, as in the situation described in (1.20), Edna has been told that she can take any of Chuck’s red vases, which Chuck considers ugly. Edna can use *lovely* to add her own comment on the red vases. But continuing this with a denial of the content of *lovely* is infelicitous.

(1.31) “Chuck said I could have one of his lovely vases. #But they are all so ugly!”

This is coherent only if we shift to a non-CI reading of *lovely*, on which its content is part of what Chuck said to Edna. (This reading is likely to imbue *lovely* with something like irony or sarcasm, given the contrasting value expressed by Edna in the continuation.) In this case, the example is irrelevant to the present discussion because it lacks a CI contributor. It is true that we call upon the maxims to understand which realization *lovely* is likely to have (CI contributor or at-issue modifier). But once the semantic translation is fixed, the maxims are not relevant. The compositional semantics does the work of determining its meaning. The other CI expressions discussed here display the same level of independence from the principles of cooperative conversation.

The distinction between conversational and conventional implicatures hinges largely on the property of deniability (cancellability). When it comes time to write a grammar, we should treat CIs as regular logical entailments. We needn't worry about sentence-external factors removing content that is usually present in different contexts. In fact, to take such information into account would be to needlessly complicate the theory of CIs. In contrast, the theory of conversational implicatures cannot escape these difficult contextual factors. On the contrary, the defining feature of such a theory will have to be a sophisticated account of how the context, the maxims, and the cooperative principle conspire to produce (and then perhaps alter or erase) this always-negotiable content.

At present, it seems fair to say that the formal theory of conversational implicatures is still quite a ways from completion (Beaver 2001:29–30). This might at first sound surprising. After all, we have a variety of compelling methods for how to calculate potential conversational implicatures and then distinguish them from actual ones; important works include Gazdar 1979a,b, Chierchia 2001, and Sauerland 2001. But all these accounts must call upon the Gricean maxims at a metagrammatical level, which is just to say that they do not succeed in bringing the maxims into the grammar itself.

This lacuna is most easily seen via a consideration of the neo-Gricean perspective developed by Chierchia (2001), who seeks to introduce conversational implicatures at the level of specific lexical items. One might think that this move reduces the distance between conversational implicatures and CIs. But in fact it does not. The deniability property still stands between them; the 'neo-Gricean' picture does not impact the design of the tree in figure 1.1.

Chierchia's (2001) starting point is the observation that scalar conversational implicatures can be embedded. I illustrate using the connective *or*, a member of the scale

$<$*or*, *and*$>$ (the stronger element is on the right).

(1.32) a. Eddie: "Mary will run the meeting or Mary will operate the projector."
b. Eddie believes that Mary will run the meeting or Mary will operate the projector.

The maxims of quantity and quality conspire to ensure that speakers always express the most informative (relevant) proposition that they have evidence for. Hence, (1.32a) is likely to conversationally implicate the falsity of the proposition that Mary will both run the meeting and operate the projector. Chierchia observes that the same implicature arises in (1.32b), though here the scalar coordinator is embedded. A global computation of conversational implicatures might wrongly predict only the weaker scalar implicature expressible as *It is false that Eddie believes that Mary will run the meeting and Mary will operate the projector.* To ensure a more local calculation, Chierchia places the scalar implicature in the lexical meaning of the determiner. In (1.33), I provide a simplified (i.e., non-type-polymorphic) version of his lexical entry for *or.*

(1.33) a. at-issue: $\lambda p \lambda q.\, p \vee q$
(*classical disjunction*)

b. conversational implicature: $\lambda p \lambda q.\, \neg(p \wedge q)$
(*classical negated conjunction*)

This looks much like the sort of meaning we have for some nodes in the CI-containing trees of later chapters. I stress, though, that the conversational-implicature dimension must be treated as formally distinct from the CI dimension studied here. At a technical level, the conversational-implicature dimension must interact with other operators:

scalar implicatures under negation (and other downward entailing operators) disappear or are radically altered, for example. The very fact that we can semantically embed the conversational implicature in (1.33) points up a distinction with CIs, which are invariant in these environments, as seen already in the initial examples in (1.1)–(1.2) and discussed more fully in later sections. The result is that they can be computed quite locally to the lexical item that triggers them.

But here is the heart of it: nothing about Chierchia's composition for *Mary will run the meeting or Mary will operate the projector* defeats the scalar implicature expressible as *It is false that Mary will run the meeting and operate the projector.* The root node for the parsetree of this sentence thus has the pair of meanings in (1.34) (ignoring tense), if we adopt the lexical entry in (1.33):

(1.34) a. at-issue:

$\mathbf{run}(\mathbf{the}(\mathbf{meeting}))(\mathbf{mary}) \vee \mathbf{operate}(\mathbf{the}(\mathbf{projector}))(\mathbf{mary})$

b. conversational implicature:

$$\neg \left(\begin{array}{c} \mathbf{run}(\mathbf{the}(\mathbf{meeting}))(\mathbf{mary}) \quad \wedge \\ \mathbf{operate}(\mathbf{the}(\mathbf{projector}))(\mathbf{mary}) \end{array} \right)$$

But the utterance might be followed in the discourse by "Hey, she'll do both!". Or it might be preceded by an agreement that if Mary does one, then she does the other. The maxims of quality and quantity would then conspire to ensure that (1.34b) disappears. Thus, the conversational-implicature dimension is a negotiable part of denotations. Even after building conversational implicatures into the compositional semantics, we still call upon the maxims to determine where they actually arise. Stepping back, we see that even if we adopt the neo-Gricean perspective, all the arguments for the distinction between the two classes of implicature hold true. The neo-Gricean perspective is just a precise, lexical method for determining where *potential* conversational implicatures lie. But their ultimate realization is something we still cannot predict without the basics of Grice's (1975) framework.

1.4.2 CIs versus at-issue entailments

The facts reviewed in section 1.4.1 leave open an analysis of CIs as at-issue entailments. Because both classes fall under the heading 'entailment', attempts to reduce the facts about CIs to at-issue meanings constitute the most pressing alternatives.

Clause (1.10d) says, in no uncertain terms, that CIs are distinct from at-issue meanings (Grice's 'what is said (in the favored sense)'). So, by stipulation, these two classes are disjoint. If this were the only point of contrast between CIs and at-issue meanings, then the distinction would arguably be a false one, perhaps simply the consequence of defining at-issue content too narrowly, or arbitrarily. This seems a fair articulation of Bach's (1999) position. Bach says that CIs are a myth, but rejects a 'one sentence, one proposition' view, offering evidence that a single sentence can express multiple nonconjoined at-issue propositions. Since Bach's descriptions implicitly appeal to a multidimensional logic (see his p. 351), it is worth seeing if Grice's definition (1.10) entails further differences between CIs and at-issue entailments.

Clause (1.10c) entails just such an additional split. A rigid interpretation of this clause (the one I adopt) means that a CI is never relativized to the beliefs of an entity other than the speaker. But at-issue content certainly is; in *Sue wrongly believes Conner got promoted*, the at-issue proposition that Conner got promoted is asserted to hold only in *Sue's* belief worlds. Thus, this embedded proposition is not speaker-oriented, and hence not classifiable as a CI contribution, by (1.10c). We can set this example alongside (1.35) to highlight the differing entailments.

(1.35) Sue wrongly believes that that jerk Conner got promoted.

This example attributes to Sue only the belief that Conner got promoted. It also involves the speaker's characterization of Conner with *that jerk*. True, Sue might also feel negatively toward Conner, thereby imparting the sense that she endorses the characterization. But this is not an entailment of (1.35). We could precede or follow the example with *Sue thinks Conner is a great guy*. However, placing (1.35) in the same context as *I think Conner is a great guy* is likely to lead to infelicity. (I refer to chapter 4 for a fuller discussion.)

The presupposition holes (negation, questioning, modalization, and conditionalization) provide even sharper judgments, with the same consequences for these meanings. All of the following carry the CI that the speaker disapproves of having to look after Shelia's dog.

(1.36)
a. I am not looking after Sheila's damn dog while she is on holiday.
b. Am I looking after Sheila's damn dog while she is on holiday?
c. I might look after Sheila's damn dog while she is on holiday.
d. If I look after Sheila's damn dog while she is on holiday, then I expect to get paid.

These observations provide initial motivation for taking seriously the claim that Grice's (1975) definition (1.10) has linguistic relevance. Establishing this claim in the face of alternatives that call upon scope-shifting mechanisms is a more difficult and involved task, one that occupies part of the argumentation in later chapters.

1.4.3 CIs versus presuppositions

Invariance under presupposition holes is consistent with an analysis of CIs as a species of presupposition. But the constructions discussed here share few properties with presuppositions; the classification seems motivated only by an attempt to cram all non-at-issue meaning into the presupposition category. Even writers not concerned directly with CIs have observed that this is inappropriate. This section mounts a multipronged attack on this reduction, using mainly supplements to motivate the claims. The arguments hold also for the other constructions reviewed above, but their slipperier content would complicate the discussion unnecessarily.

1.4.3.1 Independence of truth values

As with conversational implicatures, we can home in on the defining difference between CIs and presuppositions. Here, it is clause (1.10d), which specifies that CIs are independent of the at-issue content. In contrast, the fundamental goal of almost all presupposition logics is to create a dependency between the presuppositions and the at-issue entailments. This is the guiding intuition behind the reconstruction of presuppositions in terms of partial logics: if expression E's presuppositions are not true, then E should lack a defined value. (Karttunen and Peters (1979) might dissent from this statement. It depends on whether or not they intend their logic to model presuppositions in the usual sense.)

The exciting report in (1.37) nicely illustrates how the at-issue and CI dimensions operate independently.

(1.37) Lance Armstrong, an Arkansan, has won the 2002 Tour de France!

I know that Armstrong is a Texan; the CI is false. But I can still recover from (1.37) the information that Lance won the 2002 Tour. I need not accommodate the CI proposition to do this. In a two-dimensional semantics, the situation is easy to describe in terms of truth values. If we stick to sentences containing one at-issue value and one CI value, we have a four-valued system akin to Herzberger's (1973) logic:

(1.38)
$$\begin{array}{ll} \langle 1,1\rangle & \langle 0,1\rangle \\ \langle 1,0\rangle & \langle 0,0\rangle \end{array}$$

In our world, the extensional value of (1.37) is $\langle 1,0\rangle$. In worlds where Armstrong is neither an Arkansan nor the 2002 Tour winner, (1.37) denotes $\langle 0,0\rangle$. Neither situation should yield undefinedness for (1.37). We require both these values.

The values $\langle 1,0\rangle$ and $\langle 0,0\rangle$ are the bane of a multidimensional theory of presuppositions. These represent situations in which the presuppositions are false. One must either collapse these values to 'undefined' (Beaver 1997:956; Krahmer 1998:143), or else admit only those valuations in which presuppositions are true (van der Sandt 1988:21). One move or the other would be necessary to capture the intuition that *Ali doesn't realize her coat is on fire* is undefined if the presupposition that her coat is on fire is false.

1.4.3.2 Antibackgrounding

The dependency of at-issue meanings on their presuppositions is the most important *theoretical* divide between these meanings. The most important *pretheoretical* divide is this: CI expressions usually offer information that is not part of the common ground when they are uttered. Although it is possible for true presupposition triggers to introduce novel information, this is accompanied by a particular discourse effect, viz., accommodation. In order to understand the utterance, the hearer must adjust his knowledge so that it entails whatever the speaker has presupposed. Outside of specialized discourse conditions, it is not possible to eschew accommodation — the adjustment is thrust upon any listener who wishes to use information provided by the utterance. As Heim (1992:215, fn. 6) says, following Soames (1989:578–579), "there is no *de jure* accommodation" of a proposition $\mathbf{p}$ unless the context entails the negation of $\mathbf{p}$ already (and hence accommodation of $\mathbf{p}$ would "give rise to a communicative impass"; Soames 1989:579).

Supplements do not function in this way; their primary discourse function is to introduce new, but deemphasized material. Beaver (2001) makes this observation, and supports it with an example so lovely it is worth repeating:

(1.39) "Sweden may export synthetic wolf urine — sprayed along roads to keep elk away — to Kuwait for use against camels."[7]

Beaver observes that the proposition that wolf urine is sprayed along the roads to keep elk away is surely not part of the common ground. It is offered as new information — an aside, to be sure, but not something that the reader is expected to know already. In sum, the appositive does not express backgrounded information. We can strengthen this claim to an *antibackgrounding requirement*: in cases where the content of a supplement is part of the initial context, the result is infelicity due to redundancy, as in (1.40a).

[7] Associated Press, January 19, 1995 (cited in Beaver 2001:20, (E34)).

(1.40) Lance Armstrong survived cancer.

a. #When reporters interview Lance, a cancer survivor, he often talks about the disease.

b. And most riders know that Lance Armstrong is a cancer survivor.

With (1.40) part of the context, the use of the factive predicate *know* in (1.40b) requires no accommodation of the content of its complement. That is, the unqualified felicity of (1.40b) is contingent upon the presence of (1.40). But the same kind of backgrounding renders the appositive in (1.40a) infelicitous. As with at-issue content, we have an anti-backgrounding effect (see the partial formalization of the maxim of quantity in Groenendijk 1999:116). Neither at-issue content nor CI content should be presupposed.

This suffices to show that supplements do not meet the main pretheoretical requirements for counting as presupposed. The technical definition of 'presupposition' is much more flexible, though. Recent theories of presupposition (or, at least, recent uses of the term) somewhat weaken the strength of this argument. Steedman (2000:654) allows that "the listener rapidly and unconsciously adjusts his or her model of the domain of discourse to support the presuppositions of the speaker". If this can happen, then the difference between at-issue meanings and presuppositions is outside the bounds of detection by the usual sorts of linguistic argument. If accommodation is unconscious and freely available, then it is not distinguished from the sort of adjustments that speakers make to their models (world-views) when they accept new information. It is hard not to regard this as a complete assimilation of presuppositions to at-issue meanings. It does not accord with colloquial uses of the term 'presupposition', though.

Nonetheless, since it might be that backgrounding is not a point of contrast between

CIs and presuppositions in certain theories, I move now to some other areas of contrast. We still have deniability and presupposition plugs to turn to for support.

1.4.3.3 Deniability

Deniability (cancellability) is a hallmark of presuppositions. Green (2000) identifies this feature as one of the few things that all presupposition researchers agree upon. Green writes that “according to a wide consensus presuppositions are essentially cancelable” (p. 461). Cancellation typically arises in situations in which presuppositions conflict with the demands of the context, as in the following variation on an old example:

(1.41) Ali’s brother isn’t bold: Ali doesn’t have a brother!

If the presupposition that Ali has a brother, triggered by *Ali’s brother*, were to project, then this discourse would entail both that Ali has a brother and that she doesn’t have a brother. Felicity demands that the presupposition be filtered off. There are many theories that can obtain this result (Gazdar 1979a,b; Beaver 2001).

But as noted in section 1.4.1, CIs are never deniable in this way. The above quotation from Green 2000 occurs in the context of his argument that supplements are not presupposed. He follows up with the example in (1.42), which he calls “simply bizarre” (p. 465).

(1.42) #Snow is not white. Therefore, if, as is the case, snow is white, then grass is green.

(Green 2000:465, (26), with the judgement added)

Green's choice of examples might be regarded as unfortunate, because his *As*-parenthetical is a factive predication. We can remove this difficulty but retain the advantages of using the antecedent of a conditional:

(1.43) The press said nothing about Ames. #But if, as the press reported, Ames is a spy, then the FBI is in deep trouble.

Having stated that the press said nothing about Ames, one cannot felicitously use an *As*-parenthetical to say that the press reported that Ames is a spy, even when the clause containing this *As*-parenthetical is conditionalized.

Supplemental CI contributors are not the only ones that display this behavior. We saw above, in example (1.31), that expressive modifiers have this property as well. These cases are representative: CIs, like at-issue entailments but unlike presuppositions and conversational implicatures, cannot be cancelled.

Here again, though, we must temper this conclusion a bit. Though Green is certainly correct in his assessment that all presupposition theories attempt to come to grips with cancellability, it seems clear that this cancellability is not the same sort of thing that makes conversational implicatures so malleable. Cancellation often depends on specific intonation contours or constructions. The extreme specialization suggests an account on which presuppositions are not cancelled due to specific discourse factors broadly speaking, but rather due to specific structural configurations (being in the scope of a special kind of negation, for example; Horn 1989; Geurts 1998). Thus, I conclude on a measured note: we have some descriptive contrasts concerning the ease with which certain expected meanings are suppressed. On this score, CI expressions are significantly less yielding than presuppositions, a fact that we capture easily if the two kinds of meaning are kept apart.

1.4.3.4 Plugs don't plug them

The presupposition plugs (verbs of saying and other performatives) deliver another argument against reducing the constructions discussed here to presuppositions. Although Karttunen (1973:177) observes that "all the plugs are leaky", in the sense that they sometimes allow presuppositions to escape them, it is in general the case that a plug stops presupposition inheritance. For instance, in (1.44), the proposition that it is raining is presupposed in virtue of being expressed by the complement to *realize*. But the sentence as a whole lacks this factive presupposition.

(1.44) Ed said that Sue realized that it was raining. (Later, we found out that Ed's report was wrong. Sue can't have realized it was raining, because it wasn't.)

Karttunen is always careful to qualify this, by noting that Gricean maxims of cooperative conversation often conspire to make it seem as though an at-issue entailment or presupposition has slipped through a plug. But such content is always easily cancelled, as one would expect from a conversational implicature (Karttunen 1971; Karttunen 1973:6; Karttunen and Peters 1979:20, fn. 8). This is the utility of the continuation in (1.44), in which the speaker explicitly backs off from any hint that the complement to *realize* should be interpreted as a main clause assertion.

Here again, we find that CIs behave differently. We get a rather close minimal pair by setting (1.44) alongside the *As*-parenthetical in (1.45).

(1.45) Ed said that, as Sue reported, it is raining.

It is an entailment of (1.45) that Sue reported that it is raining, despite the appearance of the *As*-parenthetical expressing this inside the finite complement to *say*. One easy way to test the status of this proposition as an entailment is to attempt to deny it:

(1.46) Ed says that, as Sue predicted, it is raining. #But in fact Sue didn't predict rain.

The *As*-parenthetical's content is entailed even when inside a plug. This puts *As*-parenthetical content on the 'CI' branch of figure 1.1.

In the formalization of CIs in chapter 2, the unpluggability of CI content is not an extra fact. Rather, it follows directly from the basic assumption in (1.47), a slightly formal version of the intuition that CIs are comments upon an at-issue core.

(1.47) There are no types of the form $\langle \sigma, \tau \rangle$, where σ is a CI type and τ is an at-issue type.

Though a full argument must wait until the formalism is presented, the basic idea is easily put. Suppose, for contradiction, that a CI appears as part of the argument to an at-issue expression α. Then α is an at-issue term of a type that takes CI inputs to at-issue results. But (1.47) says we have no such types. No heritage function or related devise is necessary to ensure that, for each operator O, CI content is invariant under O. The type-theoretic space in which we work leaves no room for terms that violate this generalization.

1.4.4 CIs versus intonational meanings

Though current theories of intonational meaning do not provide even the basis for an analysis of the above expressions in terms of intonation, intonational meaning plays a leading role in the analysis of supplements in chapter 3. Moreover, alternative semantics for focus is a widely accepted multidimensional theory of meaning. It shows that moving beyond the first dimensional is not a suspicious or exotic move.

I do not include intonational meaning on the tree in figure 1.1, because it functions primarily as a signal for non-at-issue content, rather than as an autonomous class in itself. For example, focus meanings can impart conversational implicatures; *Joan*$_{\mathbf{F}}$ *passed* conversationally implicates that others passed, or tried to pass, simply because the focus on *Joan* invokes the set of alternative predications $\{x \mid x \text{ passed}\}$ for all contextually salient alternatives x to the individual Joan. Hearers expect, by the maxims of quantity and relevance, to find a use for the additional information.

Focus meanings can also generate presuppositions; the focus particle *only* provides an example that contrasts in important ways with CIs. In most analyses (Rooth 1992; Büring and Hartmann 2001), *only* denotes a functor that takes focus meanings and returns at-issue content. For example, the at-issue content of *Only Bob smokes* is paraphrasable as ‘No member of the set of focus alternatives to Bob smokes’. To arrive at this meaning, **only** applies to the focus meaning of *Bob* to return an at-issue quantifier. In contrast, as conceived of here, at-issue content never applies to CI content. Thus, the CI dimension must be assigned different formal properties from the focus dimension.

So there is a relationship between CIs and intonational meanings. But it is not one of subsumption. Rather, intonation is often what delivers CI content. On the analysis of supplementary relatives advocated in chapter 3, for example, the primary difference between (1.48a) and (1.48b) is the presence of a feature COMMA in the first but not the second.

(1.48) a. the crook, who has robbed hundreds of surf shops,

b. the crook who has robbed hundreds of surf shops

The feature COMMA, signalled in print by commas, dashes, or parentheses, demands that the phrase it dominates have its own intonational phrase. It also instigates a shift from at-issue meanings to CI meanings. In conjunction with the tree-admissability conditions of the CI logic, this derives the various contrasts between these two kinds of relative. This reasoning extends to a wide range of supplements and their integrated counterparts.

1.4.5 Closing remarks on kinds of meaning

This section sums up the above results, with an eye towards broader issues of how to divide up the kinds of meaning found in expressions.

A major result of the above is that one must be cautious when deploying presupposition holes to diagnose presuppositions, even when deniability is not an issue, when "the linguistic context provides no relevant information about the speaker's attitude towards" the presupposition (Beaver 2001:18). Some researchers seem to regard holes as providing both necessary and sufficient conditions for presuppositional status; Beaver (1997) writes that holes are often regarded as providing

(1.49) "an objective basis for the claim that there is a distinct presuppositional component to meaning, and a way of identifying presuppositional constructions, a linguistic test for presupposition on a methodological par with, for instance, standard linguistic constituency tests."

(Beaver 1997:945)

But CIs are another class of expressions that project beyond the holes (a fact that Beaver is aware of; see (1.39)). However, they do not display the other main criteria for presuppositions: they are not deniable; they need not (and usually cannot) be assumed by the speaker to be part of the common ground; and they invariably project beyond presupposition plugs. Thus, though the holes are useful for showing that a piece of meaning is not an at-issue entailment, further work must be done to determine where that meaning falls on the tree of meanings in figure 1.1. The holes might provide a necessary condition for presuppositionhood, but they do not provide a sufficient one. It would be a mistake to let this single factual test become definitional.

Another consideration is that one must be cautious about assigning theoretical content to the term 'implicature'. When one examines the details, it turns out that conversational and conventional implicatures differ from each other in numerous significant ways. They are perhaps the most unlike of all the kinds of meaning, sharing essentially no properties. Conversational implicatures are highly context-sensitive (and hence deniable), and inhere in no individual lexical item, emerging instead as properties of relations among propositions. CIs are not context-sensitive, not deniable, and manifest themselves *only* as properties of lexical items. Because the term 'conventional implicature' appeared first in Grice 1975, many associate CIs with pragmatic theory. But none of their main properties follows from pragmatic principles. They are narrowly grammatical entailments.

But CIs are not at-issue entailments. They project beyond presupposition plugs and have a semantic value that is independent of uncontroversial at-issue entailments.

We find CIs at (and only at) the intersection of the meanings that are entailed, speaker-oriented, and multidimensional. If we remove the entailment property, we find *some* (but not all) conversational implicatures. If we remove the multidimensionality property, we end up with at-issue entailments of various kinds. And if we remove speaker-orientation, we arrive at the class of expressions that Bach (1999) uses to support a multidimensional semantics that makes no reference to CIs. Meanings of this sort are addressed in chapter 6, but I should address here the prototypical (purported) example of CIs: the connective *but.* The above considerations are decisive: *but* does not contribute a CI. This is in line with the discussions of Chierchia and McConnell-Ginet (1990:353) and Bach (1999). The proper diagnosis seems to be Bach's: *but* determines two independent at-issue meanings.

Most telling is Bach's (1999:348) observation that, when *but* is a connective inside an indirect quotation, the entirety of its content is attributed, not to the speaker but to the subject of the verb of saying, as one would expect from an at-issue entailment. The following is Bach's (1999) example (1_{IQ}) (p. 348):

(1.50) Marv said that Shaq is huge *but* that he is agile.

Bach correctly identifies the contribution of *but* as part of what Marv said. The indirect quotation is felicitous only if Marv used *but* or an equivalent in his actual utterance — that is, only if the utterance entails the claim that being huge usually entails a lack of agility. Thus, the following discourse is infelicitous:

(1.51) Marv believes that being huge is a good indicator of agility. $^{\#}$Marv said that Shaq is huge *but* that he is agile.

The first sentence creates a context in which Marv denies the secondary (nonconjuntive) content generated by *but* in the second sentence. If the contribution of *but* were a CI, we would sense no inconsistency.

Another relevant consideration traces to a result of Barwise and Cooper (1981). They observe that *but* is the connective of choice when its two arguments are of differing monotonicity, whereas *and* is used when the conjuncts have like monotonicity. Some illustrative facts are given in (1.52). Though the examples involve *but* as a nominal connective, most analysts treat this as fundamentally the same (polymorphic) connective that connects sentences (Partee and Rooth 1983; Winter 2002).

(1.52) a. No student {but/*and} every professor attended the talk.
(*a downward monotonic quantifier and an upward monotonic one*)

b. Every student {*but/and} every professor attended the talk.
(*two upward monotonic quantifiers*)

c. No student {*but/and} no professor attended the talk.
(*two downward monotonic quantifiers*)

These facts seem also to militate against treating the differences between *but* and *and* in the CI dimension. Monotonicity properties are not conventionally implicated. They are metalogical properties of (classes of) determiner meanings. Hence, the fact that *but* and *and* alternate based on monotonicity is an indication that they too contrast in the CI dimension. A natural statement of Barwise and Cooper's (1981) generalization seems to require this kind of analysis.

This seems strong motivation for building everything about *but*'s meaning into the at-issue dimension. The only viable alternative is the presuppositional treatment suggested by Chierchia and McConnell-Ginet (1990:353), who write that "the contrastive nature of *but* appears to be backgrounded in general". The example they provide is (1.53).

(1.53) If Jim went to the store but bought nothing, we are in trouble.

They write that the “suggested contrast” — roughly, that going to the store usually entails buying something — “seems to be taken for granted” (p. 353). I endorse this characterization. But it is a long way from here to a presuppositional treatment. The property of “being taken for granted” is essential to all meanings; if I say, *Cats creep* I take for granted that we share an understanding of the meaning of *cats* and *creep*. This does not mean we are dealing with entirely presuppositional content. What is presupposed is merely the meanings of the words. One cannot infer from this to the claim that the words in question are presupposition triggers. In the case of *but*, such a classification appears not to square with examples like (1.54):

(1.54) If Jim whoozled the meezer but didn't smalunk, we are in trouble.

A speaker would infer that whoozling a meezer usually entails smalunking, despite the fact that this contrast cannot be taken for granted — the meanings of the words are unknown!

The leaves us with the task of developing a theory of multidimensional at-issue content. Chapter 6 explores the possibility of doing this using *product types*, a feature of many lambda calculi (e.g., van Benthem 1991). On this view, the translation of *but* is as in (1.55).

(1.55) *but* $\rightsquigarrow$

$$\lambda X \lambda Y \lambda x. \left\langle \begin{matrix} X(x) \wedge Y(x), \\ \mathsf{G}y[Y(y) \rightarrow \neg X(y)] \end{matrix} \right\rangle : \langle\langle \tau^a, t^a\rangle, \langle\langle \tau^a, t^a\rangle, \langle \tau^a, \langle t^a \times t^a\rangle\rangle\rangle\rangle$$

Product types are formed with the type constructor $\langle \cdot \times \cdot \rangle$. Product-typed terms are given as pairs of terms inside angled brackets.

The first term in the product type result in (1.55) is a generalized conjunction meaning. The second is a generic quantification that we can gloss as ‘For the most part, having the property denoted by Y precludes having the property denoted by X’. This is surely not the only meaning that *but* can contribute; Bach (1999:346) shows that the meaning is quite variable. But (1.55) is arguably the default interpretation. In situations in which the meanings of the arguments are vacuous or unknown, as in (1.54), this is the interpretation speakers arrive at.

In turn, we allow functors like **say** and **not** to have product-type arguments. I first provide an extensional meaning for *say*.

(1.56) *say* $\rightsquigarrow$
$$\lambda\langle p, q\rangle\lambda x.\left\langle \begin{array}{l}\textbf{say}(p)(x),\\ \textbf{say}(q)(x)\end{array}\right\rangle : \langle\langle t \times t\rangle, \langle e, \langle t \times t\rangle\rangle\rangle$$

This takes a product type consisting of two truth values as its first argument, an entity as its second. Both truth values are evaluated relative to the belief worlds of the value of the entity argument. The meaning is easily generalized.

Negation evidently functions differently, as the negative force is felt only on the first coordinate. This narrower form of application of the negation relation is easily captured:

(1.57) *not* $\rightsquigarrow$
$$\lambda\langle p, q\rangle.\,\langle\neg p, q\rangle : \langle\langle t \times t\rangle, \langle t \times t\rangle\rangle$$

Though *not* translates as a term taking propositional product types into same, the negation itself applies only to the first member. This accounts for the fact that *It’s just false that Shaq is huge but agile* does not negate the contrastive proposition that being huge generally precludes being agile.

The connective *but* is not the only functor that vindicates Bach’s (1999) claim that the 'one sentence, one proposition' motto is incorrect, but that does not fully match Grice's (1975) definition of CIs. Chapter 6 is concerned with this class of meanings. Part of the discussion's value is that it shows that a sentence expressing multiple propositions does not necessarily contain CI content. The definition in (1.10) is more articulated than that.

1.5 Chapter summary

This chapter was mostly stage setting. Section 1.2 reviewed the problematic introduction of the term 'conventional implicature' into the semantics and pragmatics literature (Grice 1975), concentrating on those aspects of the proposal that invoke speaker-orientation and independence from the at-issue semantics. Section 1.3 offered a brief introduction to the constructions that form the factual backbone of this work. Section 1.4 built the foundation for a theoretical argument that these constructions require an appeal to CIs, i.e., that they are not reducible to any of the other classes of meaning.

I assume a rich ontology of meanings — conversational implicatures, at-issue entailments, presuppositions, and CIs, all cross cut by intonation contour as a means for invoking non-at-issue content. The diversity makes more challenging the task of showing that CIs are a distinct class.

Speaker-orientation and multidimensionality are the guiding notions. The first is familiar; all main clauses are speaker-oriented in a manner made precise by the model theory of chapter 2. But multidimensionality might seem new. However, this is not really so. I noted above that alternative semantics for focus is multidimensional. When one begins to think about linguistic theory in dimensional terms, one quickly finds that the theory long ago moved beyond the idea that sentence denotations are one-tuples. Büring (1999) takes the two-dimensional framework of Rooth (1985, 1992) and adds a third dimension, the topic dimension. Dekker (2002) works with two presuppositional dimensions, showing that many of the problems with Karttunen and Peters' (1979) approach to presuppositions can be overcome in a dynamic setting. It is apparent that none of these dimensions reduces to any of the others. Linguists might end up with theories boasting more dimensions than even those of the most radical of modern physicists.

The next chapter introduces a logic that suffices as a metalanguage for a natural language semantic theory that takes seriously the CI dimension and recognizes its crucial role in the semantics of a broad range of lexical items and constructions.

Chapter 2

A logic for conventional implicatures

2.1 Introduction

This chapter develops a description logic for talking about and manipulating conventional implicature (CI) meanings. Among my goals is to integrate CIs smoothly into semantic theories that deal primarily with the at-issue dimension. The basic techniques for doing this were developed by Karttunen and Peters (1979), whose multidimensional semantic theory applies the fundamental logical insights of Herzberger (1973) to a Montagovian intensional logic translation of the sort found in Montague 1973 (PTQ). The idea that semantic translations can be pairs of lambda terms, each associated with an independent denotation, is the major innovation of Karttunen and Peters (1979), the driving force behind their success in showing "how model-theoretic methods of semantic interpretation can be extended to account for both truth-conditional and conventionally implicated meanings" (p. 3).

My approach is based on Karttunen and Peters' ideas but differs markedly in its implementation. A major conceptual change is that I replace Karttunen and Peters' rule-by-rule system with one based in type-driven translation (Klein and Sag 1985), thereby locating complexity in lexical denotations and simplifying the basic combinatoric system. A major result of this move is that we can more easily identify properties that are central to the logic. The patterns in the list of rules that Karttunen and Peters offer emerge as obvious consequences of the limited set of composition principles.

While I retain the broad outlines of the framework proposed by Klein and Sag (1985) (and in Heim and Kratzer 1998 and the work it influenced), CIs challenge their fundamental assumption that lexical items in a syntactic structure contribute exactly once to a composition (Klein and Sag 1985:171–174). I discuss this under the rubric of *resource sensitivity*. A central mode of combination in the system of Karttunen and Peters (1979) and the one developed here involves a composition scheme of this form:

(2.1)

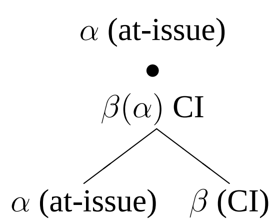

From the resource-sensitive perspective, α is consumed twice, something that is stipulated to be impossible by Klein and Sag's 'bounded closure' condition and is rendered illegitimate by the set of proof rules in Glue semantics (Dalrymple et al. 2001; Asudeh 2003) and also categorial grammar. Insights obtained in those logical systems inform my discussion of this issue, in section 2.12.

Chapter 1 touches upon a more significant divergence from Karttunen and Peters

1979. I repeat the generalization here:

(2.2) No lexical item contributes both an at-issue and a CI meaning.

For Karttunen and Peters, every expression comes with both at-issue and CI meanings (their terms are 'extensional' and 'implicature'). But the cases they discuss divide into three classes, none of which threatens (2.2). In the first case, one dimension is an identity function (as with *even*). In the second, one meaning is a conversational implicature (e.g., the name *Bill* suggests maleness). In the third, both dimensions are significant, but the lexical item in question is a presupposition trigger (e.g., *fail to*) and hence needs to be recast anyway so that it interacts properly with quantified expressions. Section 2.11 offers a fuller discussion of this issue.

The bulk of this chapter is given over to developing, step by step, a logic for CIs, which I call $\mathcal{L}_{\mathrm{CI}}$. I explain each element of $\mathcal{L}_{\mathrm{CI}}$ in depth and accompany it with its linguistic motivation. By giving the reader a feel for how it works in informal terms, I hope to make the formal apparatus as transparent and easy to work with as possible. The entire logic is presented in appendix A, with little commentary. It should be kept in mind that $\mathcal{L}_{\mathrm{CI}}$ is very much akin to the formalisms normally employed by natural language semanticists, and the models for it are entirely standard. I am careful to note, as the discussion proceeds, how results obtained in a single-dimensional framework carry over directly to this more articulated approach.

Appendix A provides the details of the notational conventions employed throughout.

2.2 Independence of truth values

Grice's (1975) definition of CIs says, in no uncertain terms, that they are separate from the at-issue dimension. It is important, he writes, that the at-issue proposition expressed by an utterance not stand or fall by the nature of its CI content. In chapter 1, I exemplified this independence using a nominal appositive (example 1.37); I repeat the example here:

(2.3) Lance Armstrong, an Arkansan, has won the 2002 Tour de France!

Cycling fans know that Armstrong is a Texan, not an Arkansan, but that he is in fact the 2002 Tour winner. The utterance nonetheless conveys information without requiring accommodation of the CI content. Generalizing, we can say that at-issue content is not dependent on the truth (or falsity) of CI content.

Nor does the converse dependency obtain. Suppose that my friend instead reports (2.4).

(2.4) Lance Armstrong, 2002's Tour winner, had never won it before 2002.

In this case, the at-issue proposition that Armstrong had not won the Tour prior to 2002 is false. Armstrong became a *four-time* winner in 2002. But I could nonetheless recover from (2.4) the proposition that Lance won in 2002. In (2.4), this is expressed in the CI dimension. Hence, the meaningfulness of the CI dimension is not dependent upon the truth (or falsity) of the at-issue dimension.

Similar observations could be made for most of the CI contributors under investigation here. For some, though, there probably are not coherent situations in which the CI could come out false. Consider a speaker's use of the expressive attributive adjective *damn* in (2.5) to express impatience with having to write a paper.

(2.5) I have to write a damn paper on fruit flies.

The at-issue entailment could be true or false. But the CI contribution of *damn* is both speaker-oriented and concerns the speaker's emotions. Since the speaker cannot, I assume, be wrong about these, the CI dimension is always true. It is true that the possibility of insincere utterances complicates the picture somewhat. But even here we run into certain difficulties in the area of expressives. Suppose that someone who is not at all racist uses a racial slur. By and large, the very act of using the slur constitutes a form of racism, whether it is sincerely used or not. The semantics developed in chapter 4 helps us make sense of this special kind of nondeniability. For now, I simply conclude that to observe the independence of the truth values of the CI and at-issue dimensions, one must study examples in which the CI dimension contributes nontrivial information about the mind–brain-external world. Facts like (2.3) and (2.4) indicate that the logic should distinguish the dimension with independent denotations: unlike logics for presuppositions, it would be a mistake to have the meaningfulness of one dimension dependent upon the truth of the other. (Chapter 1, section 1.4.3, addresses this point of contrast at a descriptive level.)

The two-dimensional semantics of Karttunen and Peters (1979), which owes much to that of Herzberger (1973), achieves the needed separation. Herzberger's logic has 'bivalence' and 'correspondence' dimensions. These are our at-issue and CI dimensions, respectively. We thus have four truth values:

(2.6)
$$\begin{array}{ll} \langle 1,1\rangle & \langle 0,1\rangle \\ \langle 1,0\rangle & \langle 0,0\rangle \end{array}$$

But sentences can contain an indefinite number of CI expressions, each contributing its own autonomous meaning. So we need to generalize the four-valued system to an n-ary valued one, by allowing that sentence denotations are (extensionally) n-ary tuples of meanings, for any finite n:

(2.7)
$$\begin{array}{llllll} 1 & 0 & & & & \\ \langle 1,1\rangle & \langle 0,1\rangle & \langle 1,0\rangle & \langle 0,0\rangle & & \\ \langle 1,1,1\rangle & \langle 0,1,1\rangle & \langle 1,0,1\rangle & \langle 1,1,0\rangle & \langle 1,0,0\rangle & \ldots \\ \langle 1,1,1,1\rangle & \ldots & & & & \\ \vdots & & & & & \end{array}$$

In short, extensional sentence meanings are members of $\{0,1\}^n$, the set of all finite tuples of truth values.

We can think of these as target denotations. To reach them, the first step is to allow certain syntactic nodes to translate as two, independently interpreted logical formulae, roughly as in (2.8).

(2.8) $$[\![\alpha : \sigma^a \bullet \beta : \tau^c]\!]^{\mathcal{M}_i,g} = \langle [\![\alpha : \sigma^a]\!]^{\mathcal{M}_i,g}, [\![\beta : \tau^c]\!]^{\mathcal{M}_i,g} \rangle$$

Here, α and β are metavariables over lambda terms, and σ^a and τ^c are metavariables over semantic types. The superscripts distinguish the types as either at-issue (superscript a) or CI (superscript c). The bullet mark, $\bullet$, is a metalogical symbol — a convenient typographic device for separating independent formulae. The interpretation brackets $[\![\cdot]\!]^{\mathcal{M}_i,g}$ are relativized to an intensional model $\mathcal{M}_i$ and a variable assignment g (as defined in (2.61) below).

Most natural language expressions for objecting to utterances target at-issue types. Karttunen and Peters (1979) discuss this at some length. Example (2.3) once again provides a helpful illustration. Suppose that I want to object to the CI proposition that

Armstrong is an Arkansan. Saying “No, that’s untrue” negates only the proposition that Armstrong won the 2002 Tour de France. It leaves the CI content untouched — the opposite of the intended effect of the negation in this context.

But there are ways to get at the CI dimension. Karttunen and Peters (1979) observe that “Well, yes, but…” is likely to indicate that the CI content is going to be disputed (p. 12). Other strategies include “Wait. I agree, but…” and even “True, but …”. The existence of these alternative strategies is a vindication of the multidimensional approach. It is impossible to make sense of a reply of the form “True, but…” in a system in which sentence meanings have just one semantic dimension.

2.3 A meaning language distinction

Chapter 1 presents a variety of contrasts between different kinds of content. The chapters following this one expand on those initial observations. Perhaps the most fundamental question that arises when devising a description logic is where to locate these distinctions among the different kinds of content. They could trace back to the syntax, the description logic, or the models — perhaps to all three. But CI expressions are syntactically heterogenous. And in a variety of cases, the models for conventionally implicated phrases are the same as those for at-issue phrases. The meaning language is thus the only viable tool for the job.

Many linguists might look askance on the rejection of a model-theoretic distinction. Given the push in much current work to directly interpret syntactic structures model-theoretically, dispensing with a description logic even in practice, a purely model-theoretic interpretation might be prima facie attractive. But it seems untenable. With regard to intersentential (discourse-level) phenomena, examples like those in (2.9) invariably pattern together, though only the first has CI content.

(2.9) a. Chuck, who killed a coworker, is in prison.
b. Chuck killed a coworker and Chuck is in prison.

These examples carry identical information. Sentence internally, this information behaves differently with respect to a number of grammatical tests (embedding, speaker-orientation, independence of semantic values). But in terms of truth conditions, (2.9a) and (2.9b) seem indistinguishable from each other. We can sharpen this intuition by looking at discourse-level phenomena. It seems that none are sensitive to the distinction between these two types of content.

Verb-phrase ellipsis and *do it*/*that* pronominalization provide a first example. Neither anaphoric device discriminates categorically between at-issue and CI meanings:

(2.10) a. Chuck killed a coworker. Sue did too.
b. Chuck, who killed a coworker, is in prison. Sue did too, but she is still walking the streets.

(2.11) a. Chuck has killed a coworker. Sue has done {that/it} too.
b. Chuck, who has killed coworker before, is in prison. Sue has done {that/it} too, but she is still walking the streets.

In the second of each pair, the antecedent for the ellipsis or proform is inside a supplementary relative clause, a CI contributor. Yet the anaphoric dependencies resolve just as they do in the first of each pair. These examples show also that CI content can serve to license additive modifiers like *too* and *also*.

Nunberg (1990) arrives at the opposite conclusion for similar examples. He writes that the content of supplements "is not actually incorporated into the text proper, and so is unavailable for external reference" (p. 105). I think that the above examples indicate that this is too strong. Admittedly it can be difficult to get an elided phrase or pronoun to find its antecedent in an appositive. But the restriction is not categorical. Rather, it seems to stem from a general preference for finding antecedents for these phenomena in the primary assertion of a preceding or nearby utterance (Frazier 2003). That processing strategy can be overridden.

We could also exploit the specific theory I offer here to help understand why it might be difficult to get anaphoric dependencies to occur between at-issue and CI expressions. In the present setting, the translation of *who killed a coworker* in (2.11b) has a CI type. But the pronominal *it* has a purely at-issue type (as does *that*). Informally, we have the following, in which g is an assignment of model-theoretic values to variables and $[\![\cdot]\!]^{\mathcal{M}_i,g}$ is an interpretation function.

(2.12) a. *it* $\rightsquigarrow f$ (*a variable with the type of at-issue predicates*)

b. *who killed a coworker* $\rightsquigarrow \lambda y.\, \exists x[\textbf{kill}(x)(y)]$
(*an expression with the type of CI predicates*)

c. For all variables f and well-formed formulae φ, if $[\![f]\!]^{\mathcal{M}_i,g} = [\![\varphi]\!]^{\mathcal{M}_i,g}$, then f and φ are expressions of the same semantic type.

The statement in (2.12c) could be either a fixed condition on admissible assignments or a descriptive generalization about how speakers are likely to understand sentences. The point is that we can achieve the distinction without resorting to the intuitively incorrect claim that *who killed a coworker* denotes one kind of function when it is a supplementary relative clause and another type of function when it is an integrated relative clause. The meaning language is sufficiently rich to provide us with conditions of the proper form, should they prove necessary.

A similar, and less fraught, test involves discourse referents. Because nothing can scope over supplementary CI content, it is quite good at establishing discourse referents. Thus, we have parallels like (2.13).

(2.13) a. Chuck spat on a coworker$_1$. $\{$She$_1$/That coworker$_1\}$ sued him and the company for $5 million.

b. They fired Chuck, who spat on a coworker$_1$. $\{$She$_1$/That coworker$_1\}$ sued him and the company for $5 million.

Here again, presupposed material patterns differently. It alone cannot establish a discourse referent, at least not without a significant amount of contextual priming and some creative inferences on the part of the hearer; well-known examples like (2.14) support this claim:

(2.14) a. Emma just got married. He's a bit grumpy.

b. Every student passed. He was thrilled by the result.

The subject pronouns in the second sentences must be deictic. For example, in (2.14b), *he* is of course unable to pick up anaphorically on the universal quantifier *every student*. But such quantifiers are commonly held to presuppose that their domain is nonempty (von Fintel 1994). Thus, the truth of the first sentence (which entails the truth of its presuppositions) guarantees the truth of the proposition that some student passed. Such existential quantifications are commonly able to establish discourse referents (it is the defining feature of most dynamic logics that they can model this). Example (2.14) shows that if the existential is merely a presupposition, it cannot function in this capacity.

A third test, also of a piece with the first, is that both CI and at-issue content can satisfy presuppositions of later sentences in a discourse:

(2.15) a. Chuck killed a coworker. And, unfortunately, his boss knows that he killed a coworker.

b. They counselled Chuck, who killed a coworker. Unfortunately, Chuck's boss knows that he killed a coworker.

The above tests are alike in being based around intersentential anaphoric dependencies — elided material, pronouns, or presuppositions, which are anaphoric on the view of van der Sandt (1992). Overall, they are an indication that we distinguish the at-issue and CI dimensions in the composition. From the point of view of a discourse as a whole, they are identical.

So far, I've called upon only supplements as support from the CI camp. This is mainly because their clausal syntax permits them to contain pronouns, indefinites, and presupposition triggers. But other CI items point to the same conclusion. Speaker-oriented adverbs, for instance, suggest a new test. Examples (2.16)–(2.17) show that their content can be queried just as at-issue content can.

(2.16) A. I was home when my parole officer called. That was unfortunate.

B. Why was that unfortunate?

(2.17) A. Unfortunately, I was home when my parole officer called.

B. Why was that unfortunate?

I should note that we have already seen one difference that relates to these examples: an objection in the form of negation is not read as an objection to the CI content. If B replied to A's utterance in (2.16) with "No", he would be read as refusing to accept that A's being home when the parole officer called was unfortunate. But in (2.17), an objection would target only the proposition that A was home when his parole officer called. The observation extends to elliptical questions. A reply of "Why?" to A's utterance in (2.16) would query A's characterization of the event as unfortunate. The same reply to A's utterance in (2.17) would query why A was present at home when his parole officer called.

These contrasts are easily made sense of in the current framework of ideas. They merely tell us that the functional types for *why* and *no* have at-issue types in their domain. Hence, we expect them to target at-issue content. The significance of (2.16)–(2.17) is that B's question *can* take the same form in both dialogues. B need not employ a special question operator for querying CI content. It seems unlikely that such special operators exist.

The upshot of these examples is that CI content is model-theoretically the same as at-issue content. Later chapters provide a wealth of evidence against a syntactic account of even a proper subset of the constructions I address here. The syntactic heterogeneity of the entire data set makes such an approach extremely unlikely to yield a fully general theory. Thus, we are lead to the conclusion that the distinction exists only at the level of the meaning language. Within this realm, one can imagine a variety of different technical implementations. The one that provides the most satisfactory formal account within the confines of present semantic theories is this: at-issue and CI content are distinguished type-theoretically. The next section explores this hypothesis in depth. Following that discussion, I head off objections that appeals to a meaning language are illegitimate. My defense is, in part, that everybody's doing it, and for good reason.

2.4 At-issue and CI types

To reduce the at-issue/CI divide to a fact about the types, the usual base step in the recursive definition of the types is divided in two: we define at-issue entities, truth values, and worlds, as well as CI entities, truth values, and worlds. The result is that we can regulate semantic composition in the meaning language, without necessarily positing model-theoretic reasons for why a given instance of functional application fails.

I provide in (2.18) the definition of the set of types for $\mathcal{L}_{\text{CI}}$, the logic employed throughout this work. The next few sections explain and justify each of its clauses.

(2.18)
a. e^a, t^a, and s^a are basic at-issue types for $\mathcal{L}_{\text{CI}}$.
b. e^c, t^c, and s^c are basic CI types for $\mathcal{L}_{\text{CI}}$.
c. If σ and τ are at-issue types for $\mathcal{L}_{\text{CI}}$, then $\langle \sigma, \tau \rangle$ is an at-issue type for $\mathcal{L}_{\text{CI}}$.
d. If σ is an at-issue type for $\mathcal{L}_{\text{CI}}$ and τ is a CI type for $\mathcal{L}_{\text{CI}}$, then $\langle \sigma, \tau \rangle$ is a CI type for $\mathcal{L}_{\text{CI}}$.
e. If σ and τ are at-issue types for $\mathcal{L}_{\text{CI}}$, then $\langle \sigma \times \tau \rangle$ is a product type for $\mathcal{L}_{\text{CI}}$, a subset of the set of at-issue types for $\mathcal{L}_{\text{CI}}$.
f. The full set of types for $\mathcal{L}_{\text{CI}}$ is the union of the at-issue and CI types for $\mathcal{L}_{\text{CI}}$.

I adopt a syntactic view of these types: rather than acting merely to index sets of denotations (Montague 1970b; Halvorsen and Ladusaw 1979), they serve as categories for lambda terms (Barendregt 1992; Reynolds 1983; Shan 2002). In essence, the types regulate semantic composition in the same way that natural language syntactic categories regulate the projection of category labels in syntactic structures. Since typing information is essential to my analysis, I always provide terms along with their types:

where α is a term and τ is a type, the expression '$\alpha : \tau$' is glossed 'the term α is of type τ' or 'the term α is in τ', in the same way that '*dog* : N^0' would naturally be read 'the natural language expression *dog* is of category N^0'.

The first two clauses define the two classes of basic types. I should note that the only basic CI type employed in this work is t^c. This might reflect something important about CIs, namely, that they are always saturated, or propositional, meanings. But it seems premature to impose this limitation at the level of the type definition. With three basic CI types, we retain a degree of flexibility that might prove useful.

Clause (2.18c) forms functional at-issue types from types in clause (2.18a). Given any two at-issue types σ and τ, $\langle\sigma,\tau\rangle$ is a functional at-issue type. In theories in which the CI dimension is not a concern, this is the only clause for functional types. In such theories, the entire type definition might be given with (2.18a) and (2.18c) (along with a closure condition like (2.18f)).

The next clause is the main innovation of the $\mathcal{L}_{\mathrm{CI}}$ type theory. It defines functional CI types. As with at-issue types, there is just one possibility: for any at-issue type σ and any CI type τ, we have the functional CI type $\langle\sigma,\tau\rangle$. However, with τ a CI type and σ an at-issue type, $\langle\tau,\sigma\rangle$ is not a well-formed type of any kind.

The final substantive clause, (2.18e), defines product types. Expressions of this type are interpreted as ordered pairs of meanings. The product types form a subset of the at-issue, because all their subtypes are required to be at-issue.

In sum, (2.18a), (2.18c), and (2.18e) constitute the type-logical space for at-issue meanings, whereas (2.18b) and (2.18d) constitute the type-logical space for CI meanings. It is somewhat like having two type definitions, except that the clause for forming functional CI types employs both at-issue and CI types. This is how the two classes of meaning come to interact.

In general, I do not employ abbreviatory devices when it comes to type specifications. However, the reader might find it useful to keep in mind that we *could* abbreviate the superscript marking considerably. For instance, we could abbreviate $\langle e^a, t^a \rangle$ as $\langle e, t \rangle^a$, with the superscript indicating what the type is. With the information that the type is at-issue, we know that all of its subtypes are also at-issue. Similarly, $\langle e, t \rangle^c$ is a potential abbreviation for $\langle e^a, t^c \rangle$. The shape of the inductive procedure is fairly evident; here are the details:

(2.19) Let x serve as a variable over $\{e, t, s\}$, and let σ and τ serve as variables over well-formed types with their superscripts stripped off. The type-superscript abbreviator $\leftrightsquigarrow$ is defined as follows:

$$\begin{array}{rcl} x^a & \leftrightsquigarrow & x^a \\ x^c & \leftrightsquigarrow & x^c \\ \langle \sigma^a, \tau^a \rangle & \leftrightsquigarrow & \langle \sigma, \tau \rangle^a \\ \langle \sigma^a, \tau^c \rangle & \leftrightsquigarrow & \langle \sigma, \tau \rangle^c \end{array}$$

I exploit this kind of abbreviation only when using metavariables to define overarching conditions, as in the definitions of the meaningful expressions and the tree-admissibility conditions. In such contexts, σ^a is a metavariable over types whose subtypes are all at-issue, and σ^c is a metavariable over types whose subtypes are all at-issue except the last one, which is a CI type. When giving actual analyses — when using the types themselves — I've chosen not to take advantage of this typographic simplification because it can hinder comprehension of the way individual composition schemes work. I provide the details in (2.19) mainly because they help to bring out the tight logical structure of the type definition.

A partial definition of the set of well-formed terms is given in (2.20). The definition is partial in that it provides only the clauses of immediate importance to the general combinatorics; the clauses for the connectives and quantifiers are given in the definition in appendix A, along with the rest of $\mathcal{L}_{\mathrm{CI}}$.

(2.20) Let ME_τ denote the set of all meaningful expressions of type τ for $\mathcal{L}_{\mathrm{CI}}$.

i. If c is a constant of type τ, then $c \in ME_\tau$.
ii. If x is a variable of type τ, then $x \in ME_\tau$.
iii. If $\alpha \in ME_{\langle \sigma^a, \tau^a \rangle}$ and $\beta \in ME_{\sigma^a}$, then $(\alpha(\beta)) \in ME_{\tau^a}$.
If $\alpha \in ME_{\langle \sigma^a, \tau^c \rangle}$ and $\beta \in ME_{\sigma^a}$, then $(\alpha(\beta)) \in ME_{\tau^c}$.
iv. If $\alpha \in ME_{\tau^a}$ and x is a variable in ME_{σ^a}, then $(\lambda x.\, \alpha) \in ME_{\langle \sigma^a, \tau^a \rangle}$.
If $\alpha \in ME_{\tau^c}$ and x is a variable in ME_{σ^a}, then $(\lambda x.\, \alpha) \in ME_{\langle \sigma^a, \tau^c \rangle}$.
v. If $\alpha \in ME_{\sigma^a}$ and $\beta \in ME_{\tau^a}$, then $\langle \alpha, \beta \rangle \in ME_{\langle \sigma^a \times \tau^a \rangle}$.
vi. The full set ME of meaningful expressions is the union of all the sets ME_τ for all types τ.

I almost always drop outermost parentheses. I assume also that application associates to the left. Thus, $((\alpha(\beta))(\gamma))$ abbreviates to $(\alpha(\beta))(\gamma)$ by the convention that drops outermost parentheses. We can abbreviate further to $\alpha(\beta)(\gamma)$ by the convention that associates application to the left.

The local trees in (2.21) provide a sense for how this articulated type definition works for simple cases.

(2.21) a. $(\lambda x.\, \textbf{cyclist}(x))(\textbf{lance}) : t^a$

$\lambda x.\, \textbf{cyclist}(x) : \langle e^a, t^a \rangle$ $\quad$ $\textbf{lance} : e^a$

b. undefined

$\lambda x.\,\mathbf{cyclist}(x) : \langle e^a, t^a \rangle$ $\quad$ **erik** : e^c

The first tree has a legitimate root-node translation, by functional application of its daughters (clause (2.20iii)). It is decorated with a term of type t^a, one that converts to **cyclist**(**lance**). But the labels on the daughter nodes in the second tree cannot combine to yield a well-formed label for their mother. Because $\lambda x.\,\mathbf{cyclist}(x) : \langle e^a, t^a \rangle$ requires an at-issue entity as its first argument, the individual constant **erik**, which is a CI entity, is not a suitable argument. In the models, **lance** and **erik** denote individuals, and $\lambda x.\,\mathbf{cyclist}(x) : \langle e^a, t^a \rangle$ denotes a function from individuals to truth values. But the meaning language blocks (2.21b).

2.5 Linguistic motivation for the limited set of types

The guiding idea behind the set of types defined in (2.18) is that CIs bear an asymmetric relationship to at-issue meanings: they apply to at-issue meanings to produce CI meanings. We do not have at-issue meanings applying to CI meanings, regardless of the resulting type. The asymmetry encodes, as transparently as possible, the idea that CIs are comments upon an asserted core.

The clauses that achieve this result are the ones defining functional types. I repeat them here with their original numbering:

(2.18) c. If σ and τ are at-issue types for $\mathcal{L}_{\mathrm{CI}}$, then $\langle \sigma, \tau \rangle$ is an at-issue type for $\mathcal{L}_{\mathrm{CI}}$.

d. If σ is an at-issue type for $\mathcal{L}_{\mathrm{CI}}$ and τ is a CI type for $\mathcal{L}_{\mathrm{CI}}$, then $\langle \sigma, \tau \rangle$ is a CI type for $\mathcal{L}_{\mathrm{CI}}$.

These clauses have glosses that capture directly my earlier descriptive statements about how the combinatorics work; I summarize these in (2.22).

(2.22) a. At-issue meanings apply to at-issue meanings to produce at-issue meanings.

b. CI meanings apply to at-issue meanings to produce CI meanings.

The first possibility is the bread and butter of formal semantics. The second reflects the intuition that CIs borrow from the at-issue dimension. We expect not to find cases where the at-issue meaning applies to a CI meaning, as this would undermine the characterization of CIs as peripheral, nonintrusive commentary.

2.5.1 At-issue never applies to CI

The primary advantage that the types bring is that they regulate composition via the set of well-formed lambda terms. The first term in (2.23) is well formed, whereas the second is not.

(2.23) a.

$$\lambda x.\, \textbf{believe}(\textbf{cyclist}(\textbf{lance}))(x) : \langle e^a, t^a \rangle$$

$$\lambda p \lambda x.\, \textbf{believe}(p)(x) : \langle t^a, \langle e^a, t^a \rangle\rangle \qquad \textbf{cyclist}(\textbf{lance}) : t^a$$

b.

$$\text{undefined}$$

$$\lambda p \lambda x.\, \textbf{believe}(p)(x) : \langle t^a, \langle e^a, t^a \rangle\rangle \qquad \textbf{cyclist}(\textbf{lance}) : t^c$$

Tree (2.23b) suffers a type mismatch: the functor is of a type that requires an at-issue truth value; the argument is of type t^c. The two cannot combine directly to yield a term of the logic.

The type-mismatch in (2.23b) represents a desirable gap in the type-logical space defined in (2.18). The functor is an appropriate extensional meaning for a propositional attitude verb. The CI proposition in its complement cannot serve as its argument. Moreover, we cannot have verb meanings that take CI meanings as their arguments. Such a verb would have to be of type $\langle t^c, \langle e^a, t^a \rangle\rangle$, that is, it would have to have an initial member of CI type but itself be of at-issue type. Such types are not in the set defined in (2.18). Analogous reasoning applies to operators like negation, conditionalization, modals, and the like. This is the basis for the explanation for why CIs never form part of the argument to these operators. For instance, we saw in chapter 1 that expressive attributive adjectives like *damn* are not part of the argument to higher functors. A simple example illustrating this behavior:

(2.24) a. Bush says the damn Republicans deserve public support.

b. Clinton: The damn Republicans want the bill passed.

Bush: #Clinton says the damn Republicans want the bill passed.

Example (2.24a) is a complete and accurate report of Bush's utterance "The Republicans deserve public support". The pair of utterances in (2.24b) sharpens the intuition, as it is extremely hard to imagine that a die-hard Republican would report Clinton's utterance in this way; even those with a limited grasp of the language would recognize that *damn* will be attributed to the speaker unless given a special intonation contour indicating that it is intended as a quotative utterance (the usual signal would be heavy emphasis on the expressive attributive adjective).

The explanation for this filtering is simple. The verb *say* has a translation like that represented in (2.23), which I intensionalize in (2.25).

(2.25) $\quad$ *say* $\rightsquigarrow \lambda \mathbf{p} \lambda x \lambda w.\, \mathbf{say}_w(\mathbf{p})(x) : \langle \langle s^a, t^a \rangle, \langle e^a, \langle s^a, t^a \rangle \rangle \rangle$

In light of the discussion in section 1.4.5 and in chapter 6 below, we probably want to generalize this meaning so that it can take product-type arguments of any level of complexity. In order to do this relatively perspicuously, I first exploit the commutativity of products to abbreviate, for example, $\langle \sigma \times \langle \tau \times \upsilon \rangle \rangle$ as $\langle \sigma \times \tau \times \upsilon \rangle$. Then we can abbreviate a run of identical product types as follows:

(2.26) $\quad \langle \sigma_1 \times \cdots \times \sigma_n \rangle \leftrightsquigarrow \overrightarrow{\langle \sigma \rangle}_n$

It is now possible to give a generalized meaning for *say*, one that permits it to apply to arguments that are tuples of at-issue propositions. I provide such a meaning in (2.27), which abbreviates an infinite set of terms of $\mathcal{L}_{\mathrm{CI}}$.

(2.27) $\quad$ *say* $\rightsquigarrow \lambda \langle \mathbf{p}_1, \ldots, \mathbf{p}_n \rangle \lambda x. \left\langle \begin{array}{c} \lambda w.\, \mathbf{say}_w(\mathbf{p}_1)(x), \\ \vdots \\ \lambda w.\, \mathbf{say}_w(\mathbf{p}_n)(x) \end{array} \right\rangle : \langle \overrightarrow{\langle s^a, t^a \rangle}_n, \langle e^a, \overrightarrow{\langle s^a, t^a \rangle}_n \rangle \rangle$

The type does not change; the tuple of propositions that is its first argument consists entirely of at-issue propositional terms. Thus, even if *say*'s clausal complement had the translation in (2.28), **say** would apply only to the at-issue proposition.

(2.28) $$\begin{array}{c} \lambda w.\, \mathbf{pass}_w(\mathbf{the}_w(\mathbf{bill}_w))(\mathbf{the}_w(\mathbf{republicans}_w)) : \langle s^a, t^a \rangle \\ \bullet \\ \lambda w'.\, \mathbf{bad}_{w'}(\mathbf{republicans}_{w'}) : \langle s^a, t^c \rangle \end{array}$$

I stress that the clausal complement in this case does not in fact have the content of *damn* on its root node; section 2.6.7 develops a theory of composition that does not pass CI propositions up in the way that this meaning suggests.

The arguments from negation and conditionalization are the same. Both involve manipulation of at-issue meanings alone.

2.5.2 CI never applies to CI

A lacuna in the combinatorics is that CI meanings don't apply to other CI meanings. Semantic analysis hinges upon at-issue meanings applying to same. One might have expected the CI dimension to work in parallel, an independent combinatoric scheme. But this seems not to be so. From the perspective of chapter 1, this is intuitively right: CI are comments upon the at-issue dimension. They are not comments upon themselves.

I've found a few opportunities to support this claim with ungrammatical or anomalous sentences, though it must be kept in mind that they can only be suggestive: they do not show that we couldn't have such composition, only that language does not avail itself of the chance in the cases at hand.

A first example is found in Kratzer (1999), whose focus is primarily the German discourse particle *ja*. When *ja* appears in a clause C, it contributes the secondary proposition that the hearer is probably aware of the content of C. Thus, one has examples like (2.29).

(2.29)	Dies	sind	ja	schlechte	Zeiten.
	these	*are*	*JA*	*bad*	*times*

'These are, as you are probably aware, bad times.'

Kratzer 1999 is a commentary on Kaplan 1999, which includes a lengthy discussion of epithets. Kratzer observes that these do not interact semantically with *ja*; of the example in (2.30), she writes that "The epithet *das Arschloch* is not in the scope of *ja*" (p. 4).

(2.30)	Sie	haben	ja	den Webster	– das Arschloch	—	endlich	gefeuert.
	they	*have*	*JA*	*the Webster*	*the asshole*		*finally*	*fired*

'They JA finally fired Webster, the asshole.' (Kratzer 1999)

The hearer is not presumed by the speaker to have prior knowledge that Webster is an asshole or is believed by the speaker to be an asshole. Only the fact that they finally fired him forms the argument to *ja*. Assuming that the contributions of the epithet and *ja* are both CI-based, we have an argument against allowing CI meanings to apply to other CI meanings.

It's possible to construct similar examples in English. We have seen that adverbs like *amazingly* can have a CI-based semantics, taking at-issue propositions into CI propositions. Suppose they could also take CI propositions into same. Then (2.31a) would have a semantic parse in which the adverb modified the CI proposition expressed by the nominal appositive:

(2.31) a. *They replied to amazingly Lance, a four-time Tour winner.
b. They replied to Lance, amazingly a four-time Tour winner.

When the adverb is inside the NA, as in (2.31b), it can take the property denoted by *a four-time Tour winner* as its argument to return a modified property, which they shifts to become CI content. (On the analysis in chapter 3, the shift is engendered by the comma intonation itself.) But *amazingly*'s meaning can also take propositions into propositions. The CI content of the object *Lance, a four-time Tour winner* is the proposition that Lance is a four-time Tour winner. And yet the ungrammaticality of (2.31a) indicates that this cannot be modified by *amazingly*.

A second example, this one requiring a more intricate construction: it is common for expressive attributive adjectives to take their immediate clause's denotation as their semantic arguments. The most likely reading of (2.32) involves this kind of composition.

(2.32) I have to mow the fucking lawn.

The speaker probably bears no ill-will towards lawns, or his lawn. Rather, the proposition that he must mow the lawn is what he seeks to disparage. Thus, we have a composition scheme like (2.33). (I defer to chapter 4 a discussion of how **fucking** ends up as a clausal modifier.)

(2.33)

$$\begin{array}{c}
\textbf{must}(\textbf{mow}(\textbf{the}(\textbf{lawn}))(\textbf{the-speaker})) : \langle s^a, t^a \rangle \\
\bullet \\
\textbf{fucking}(\textbf{must}(\textbf{mow}(\textbf{the}(\textbf{lawn}))(\textbf{the-speaker}))) : \langle s^a, t^c \rangle \\
\begin{array}{cc}
\textbf{fucking} : & \textbf{must}(\textbf{mow}(\textbf{the}(\textbf{lawn}))(\textbf{the-speaker})) : \\
\langle \langle s^a, t^a \rangle, \langle s^a, t^c \rangle \rangle & \langle s^a, t^a \rangle
\end{array}
\end{array}$$

If we had functors from CIs into same, we would expect them to be able to take the CI on this root node as their arguments. So, for example, *As*-parentheticals and supplementary relatives should be able to target the CI here. But this is not a possible analysis:

(2.34) I have to mow the fucking lawn, as my Dad said.

a. *As*-parenthetical $=$ my Dad said I have to mow the lawn

b. *As*-parenthetical $\neq$ my Dad said I disapprove of having to mow the lawn

c. *As*-parenthetical $\neq$ my Dad said I have to mow the fucking lawn

(2.35) I have to mow the fucking lawn, which is reasonable if you ask me.

a. supplementary relative $=$ that I have to mow the lawn is reasonable if you ask me

b. supplementary relative $\neq$ that I disapprove of having to mow the lawn is reasonable if you ask me

Both of the impossible readings are reasonable in terms of their information content. But these sentences cannot convey said content.

The set of types in (2.18) is designed to rule out the above sort of composition. The semantic translations required to deliver the ungrammatical examples in this section are not members of the set specified as the basis for defining and organizing lexical meanings.

2.6 Modes of combination

The hallmark of type-driven translation is that the semantic value of a syntactic node u is determined by the semantics of u's daughters and *functional application*, the interpretation of the axiom of the lambda calculus called β-conversion. By the nature of the theory of types, there is always exactly one legitimate (i.e., defined) semantic value for u — a deterministic system. The type-driven system for CIs should remain deterministic. The job of this section is to formulate general type-driven composition rules that achieve this.

I state the semantic combinatoric rules as tree-admissibility conditions, where the trees in question are *semantic parsetrees* as defined in (2.36).

(2.36) A *semantic parsetree* is a structure $\mathcal{T} = (T, \mathbf{D}, V)$, where

a. $T = \{u_1, u_2, \ldots\}$ is a set of nodes.

b. $\mathbf{D}$ is an irreflexive, intransitive binary relation on T; it is defined so that, for all $u \in T$, there is at most one u' such that $\mathbf{D}(u', u)$ and at most two distinct nodes u', u'' such that $\mathbf{D}(u, u')$ and $\mathbf{D}(u, u'')$.

c. $\mathbf{D}^*$, the reflexive, transitive closure of $\mathbf{D}$, is acyclic.

d. There is a unique $r \in T$, the root: there is no $u \in T$ such that $\mathbf{D}(u, r)$.

e. V is a valuation function, taking formulae of $\mathcal{L}_{\text{CI}}$ to sets of nodes in T, according to tree-admissibility conditions (2.38), (2.40), (2.42), (2.45), and (2.46).

Thus, the structures are connected, rooted, acyclic graphs. The branching factor for each node is at most 2, and each node has at most one mother.

We can view the logic $\mathcal{L}_{\text{CI}}$ as a specification of the parsetrees determined by (2.36). I leave open the metalogical (metagrammatical) interpretation of these trees. They can be viewed in various ways — as proof rules, tree-generation procedures, etc. I specify only that the ordering of terminal elements is irrelevant. I usually order the leaves according to their linear ordering in the syntax, but this is purely for convenience.

2.6.1 One node structures

The first condition mentioned in the definition of the valuation function, (2.36e), is given here:

(2.37) $\alpha : \sigma$ (where α is a meaningful expression of $\mathcal{L}_{\text{CI}}$)

This condition just licenses single-node trees decorated with terms drawn from the set defined in (2.20).

2.6.2 At-issue functional application

I provide here the usual clause for functional application of sisters. The definition is identical in content to those of Klein and Sag (1985:171), Heim and Kratzer (1998:44) and much other work.

(2.38) **at-issue application**

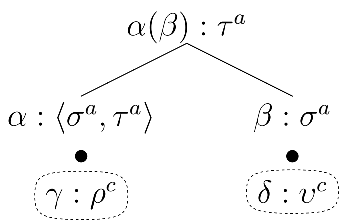

In essence, this is just clause (2.20iii) stated over semantic parsetrees. Here and in the statement of the other rules, I indicate optional material inside dotted lines. Such material is always CI content. It is separated graphically from the at-issue term above it by a bullet, $\bullet$, a metalogical device for separating independent lambda terms. The motivation for the optional material is that we must allow that there might be CI content hanging around. The rule for parsetree interpretation, (2.50), ensures that such material forms part of the overall interpretation. But it is not relevant to the local calculations that these rules determine. Thus, for example, we license both of the following trees using (2.38):

(2.39) a.

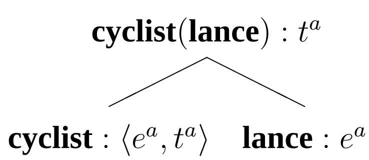

b.

$$\begin{array}{ccc} & \mathbf{cyclist}(\mathbf{lance}) : t^a & \\ \mathbf{cyclist} : \langle e^a, t^a \rangle & & \mathbf{lance} : e^a \\ & & \bullet \\ & & \mathbf{training}(\mathbf{lance}) : t^c \end{array}$$

In the second, we have a propositional CI term on one of the daughters. It is ignored in the determination of the label on the mother.

2.6.3 At-issue intersection

I offer a second method for combining at-issue meanings. The condition is called **at-issue intersection**. It is defined for all and only like-typed at-issue terms ending in t^a.

(2.40) **at-issue intersection**

$$\begin{array}{ccc} & \lambda X.\, \alpha(X) \wedge \beta(X) : \langle \sigma^a, t^a \rangle & \\ \alpha : \langle \sigma^a, t^a \rangle & & \beta : \langle \sigma^a, t^a \rangle \\ \bullet & & \bullet \\ \gamma : \rho^c & & \delta : \upsilon^c \end{array}$$

There are alternatives to this rule. It is here merely as a tool for deriving meanings for some modification structures. It plays only a supporting role in the proposals of this dissertation. It is here merely to keep the at-issue semantics running smoothly in the background. The **at-issue application** rule above could do its work, and in fact must do its work for nonintersective adjectives like *fake* and *ungrammatical*: to get the meaning of *ungrammatical* sentence, we do not intersect the ungrammatical things with the sentences, as their intersection is, by definition, empty. Rather, we must apply the meaning of *ungrammatical* to that of *sentence*, a scheme we can assure via a careful assignment of meanings to these words.

2.6.4 CI application

We come now to the central tree-admissibility condition of the CI logic $\mathcal{L}_{\mathrm{CI}}$.

It is a feature of the logic for CIs of Karttunen and Peters (1979) that a CI meaning always applies to an at-issue meaning to produce a CI meaning. This is not something they comment upon, but their grammar fragment is based in such types. The basic combination scheme underlies, for example, my analysis of speaker-oriented adverbs like *fortunately* in (2.41).

(2.41) a. Fortunately, Beck survived.

b.

$\lambda w.\, \textbf{survive}_w(\textbf{beck}) : \langle s^a, t^a \rangle$

$\bullet$

$\textbf{fortunately}(\lambda w.\, \textbf{survive}_w(\textbf{beck})) : \langle s^a, t^c \rangle$

$\textbf{fortunately} : \langle\langle s^a, t^a \rangle, \langle s^a, t^c \rangle\rangle$ $\qquad$ $\lambda w.\, \textbf{survive}_w(\textbf{beck}) : \langle s^a, t^a \rangle$

Note how the at-issue term $\lambda w.\, \textbf{survive}_w(\textbf{beck})$ is both passed on to the mother node and part of the argument to the CI adverb, with the result of such application passed on to the mother as well. The job now is to extract from (2.41) a general composition rule. I do this in (2.42).

(2.42) **CI application**

$$\beta : \sigma^a$$
$$\bullet$$
$$\alpha(\beta) : \tau^c$$
$$\alpha : \langle \sigma^a, \tau^c \rangle \qquad \beta : \sigma^a$$
$$\bullet \qquad\qquad \bullet$$
$$\gamma : \rho^c \qquad \delta : \upsilon^c$$

An important observation is that this rule functions to ensure that the at-issue dimension is always insensitive to the presence of adjoined CI operators. In other words, for any tree $\mathcal{T}$, the at-issue semantic content of $\mathcal{T}$ is the same as the tree $\mathcal{T}'$ gotten from $\mathcal{T}$ by pruning all nodes dominating items with a CI semantics (i.e., translating as a term of type σ^c). The graphic in (2.43) helps convey the intuitive content of this observation about $\mathcal{L}_{\text{CI}}$'s parsetrees.

(2.43)

$$\textbf{lance} : e^a$$
$$\bullet$$
$$\textbf{cyclist}(\textbf{lance}) : t^c$$
$$\textbf{lance} : e^a \qquad \textbf{cyclist} : \langle e^a, t^c \rangle$$

$=$

$$\textbf{lance} : e^a$$
$$\textbf{lance} : e^a$$

2.6.5 Isolated CIs

It is sometimes the case that CI expressions do not interact with the at-issue material around them in a way that is representable in terms of function application. I offer two such examples in (2.44); similar constructions are discussed at various points in the chapters to come.

(2.44) a. Luke — and you'll never believe this — ate fifty eggs

b. That's fantastic fucking news!

In the niched conjunction in (2.44a), the supplementary expression is apparently saturated. The meaning of the pronoun is identical to that of the main clause predication, but this is probably best achieved by ensuring that *this* and *Luke ate fifty eggs* pick out identical objects. Thus, the meaning of the supplement is the proposition that the addressee will never believe that Luke ate fifty eggs. Similarly, it seems wrong to say that the contribution of *fucking* in (2.44b) involves anything more integrated than the speaker's expression of a particular emotion.

What we seek for all three cases is a way to allow the CI material to remain completely separate from the at-issue content. The rule of **isolated CIs** licenses the requisite subtrees.

(2.45) **isolated CIs**

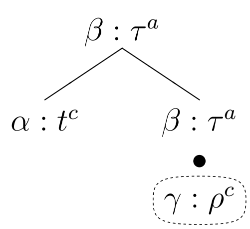

The most noteworthy feature of this rule is that it works only for adjoined CI content that is saturated. Nothing else can contribute a totally isolated meaning in this way.

2.6.6 Features

I posit a special rule, **feature semantics**, for representing the semantic contribution of certain syntactic features.

(2.46) **feature semantics**

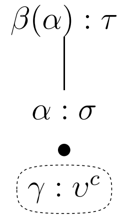

(where β is a designated *feature term* of type $\langle \sigma, \tau \rangle$)

This rule allows us to introduce features without requiring that they be terminal elements in the syntax or the semantics. The most important appeal to such a rule comes in chapter 3, in which I propose that many of supplemental expressions are distinguished by a feature COMMA, the semantics of which simply switches a limited class of expressions from at-issue to CI content. The following appositive structure is typical:

(2.47) Lance, a cyclist,

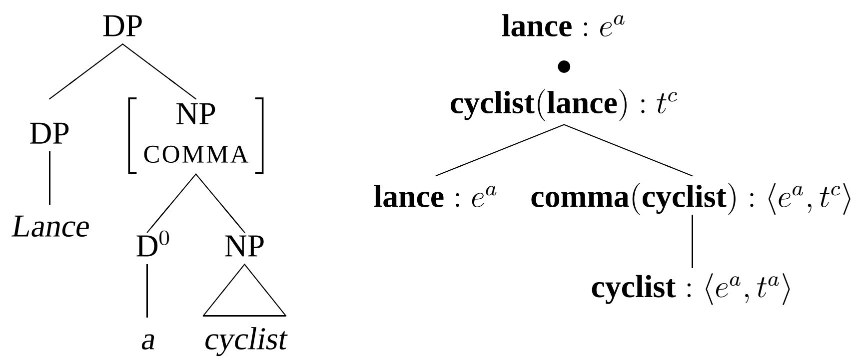

The features introduce a slight departure in the semantics from the shape of the syntactic tree: the features in the syntactic structure each have their own unary branching node in the semantics. Here is a general mapping principle:

(2.48)

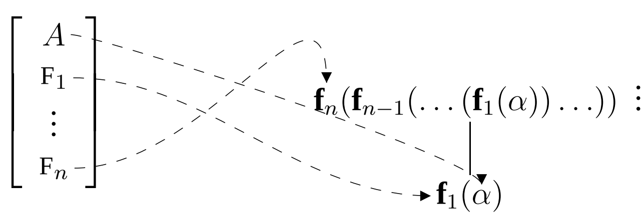

This looks less like a divergence of form between the syntax and the semantics if one views nodes heavily laden with features as relational structures in their own right. This idea was first advanced by Blackburn et al. (1993) for Generalized Phrase Structure Grammar, in which individual nodes are profitably thought of as possessing their own internal structure (see also Blackburn and Meyer-Viol 1997). On the simplest of such treatments, the nodes would have a unary branching structure. If that were the case, then **feature semantics** would not introduce deviations of form at all.

2.6.7 Parsetree interpretation

It is commonly the case that a single sentence contains phrases that express a number of different CI propositions. These are often syntactically embedded, and I've so far made no provision for what to do with them to ensure that they are part of the overall meaning of the structure. A pointed (though clunky) example:

(2.49) Dave, the boss, bribed Sue, an OSHA representative.

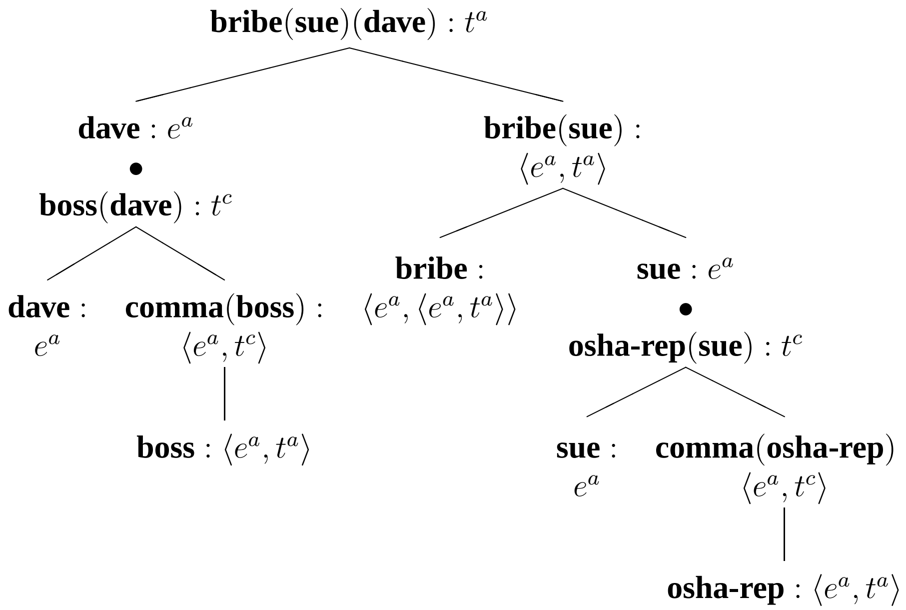

Both subject and object contain CI propositions. I see a variety of linguistically equivalent options for ensuring that these form part of the denotation of (2.49). One option would be to pass all terms of type t^c up the tree unmodified. The result would be a root node decorated with a tuple of lambda terms: one at-issue term, and n CI terms, for $n \geq 0$. This option is not the best one, because it is in effect a heritage function without any notable properties. By ensuring that CIs do not interact with other operators, we guarantee that they never end up with narrow scope. It seems a redundancy to pass them up the tree as though they were in danger of falling in the scope of something.

A second option: we could define a CI store. Any expression of type t^c would be added to the store. The interpretation of a sentence would be the interpretation of its root node, plus the interpretation of all terms in the CI store. The result would again be that sentences denote tuples of meanings, as with the heritage function solution.

I favor a third option, though: leave the CIs where they are, but interpret the entire parsetree. In other words, rather than reducing semantics to the function named on the root node, we interpret structured objects. The following definition achieves the effects of the heritage function and CI store options without the addition of new devices:

(2.50) **parsetree interpretation**

Let $\mathcal{T}$ be a semantic parsetree with the at-issue term $\alpha : \sigma^a$ on its root node, and distinct terms $\beta_1 : \langle s^a, t^c \rangle, \ldots, \beta_n : \langle s^a, t^c \rangle$ on nodes in it (extensionally, $\beta_1 : t^c, \ldots, \beta_n : t^c$). Then the interpretation of $\mathcal{T}$ is the tuple

$$\left\langle [\![\alpha : \sigma^a]\!]^{\mathcal{M}_i,g}, [\![\beta_1 : \langle s^a, t^c \rangle]\!]^{\mathcal{M}_i,g}, \ldots, [\![\beta_n : \langle s^a, t^c \rangle]\!]^{\mathcal{M}_i,g} \right\rangle$$

where $[\![\cdot]\!]^{\mathcal{M}_i,g}$ is the interpretation function, taking formulae of the meaning language to the interpreted structure $\mathcal{M}_i$, relative to a variable assignment g.

With this definition, the interpretation of (2.49) is determined by the at-issue term on its root node as well as the terms **osha-rep**(**sue**) : t^c and **boss**(**dave**) : t^c. This is the desired result. It assigns all CI content the same semantic force as a main clause assertion, simply in virtue of the fact that the root node and any CI content at or below it are all interpreted in exactly the same fashion. In effect, **parsetree interpretation** gives propositional CI content a free ride to the root node without performing any syntactic manipulation of the parsetrees or enriching them with a CI store.

2.6.8 In sum

The table in (2.51) summarizes the modes of combination allowed by the above rules and the type theory. A superscript $func$ indicates the functor, which applies to the argument meaning, superscripted with arg.

(2.51)

	at-issue^{arg}	CI^{arg}
at-issue^{func}	at-issue	$*$
CI^{func}	CI • at-issue	$*$

The asterisks indicate unattested meanings. The lower left quadrant is far and away the most important for the present study. The case where an at-issue meaning applies to an at-issue one, in the upper left corner, is just basic compositional semantics. While it plays a significant role in all analyses offered here, my interest in it is limited to making sure that the logic I offer does not disrupt or demand modifications to this kind of combination. As one can see from the way the logic is designed, this is in fact the case.

2.7 Remarks on appeals to a meaning language

A semantic translation language is commonly viewed as dispensable, at least in principle. Very often, the interpretation function $[\![\cdot]\!]^{\mathcal{M}}$ works as in (2.53): it has its domain in natural language expressions and its range in the model $\mathcal{M}$. An indirect interpretation system that is much like the one employed in this work is represented in (2.52). It adds a middle step to direct interpretation.

(2.52) *cyclist* $\rightsquigarrow$ $[\![\textbf{cyclist}]\!]^{\mathcal{M}} =$ that function f such that $f(x) = 1$ if x is a cyclist, else $f(x) = 0$

(2.53) $[\![\textit{cyclist}]\!]^{\mathcal{M}} =$ that function f such that $f(x) = 1$ if x is a cyclist, else $f(x) = 0$

We owe both schemes to Montague. The direct translation method dates to Montague 1970a (‘English as a formal language’). Indirect translation is the hallmark of Montague (1973) (PTQ), in which Montague writes “We could (as in Montague (1970a)) introduce the semantics of our fragment directly; but it is probably more perspicuous to proceed indirectly” (p. 23; ‘Montague (1970a)’ is ‘English as a formal language’). If both the translation relation $\rightsquigarrow$ and the interpretation function $[\![\cdot]\!]^{\mathcal{M}}$ are homomorphisms, then their composition is also a homomorphism, in principle eliminable without any substantive change in the theory.

But in the CI logic employed here, translation is a nontrivial step, as it is the locus of the at-issue/CI distinction. If we interpreted natural language expressions directly, we would be left with only two options, both undesirable: we could locate the distinction in the natural language syntax or in the models. The syntactic view is initially unpromising given the semantic nature of the definition of CIs, and it receives sustained criticism in chapter 5. The model-theoretic option is to locate the differences in the models. But section 2.3 shows that intersentential phenomena treat the two classes of meaning as model-theoretically identical. Moreover, it is hard to see what the distinctions could be. Assume that we want the term **cyclist** to have the meaning in (2.54).

(2.54) $[\![\textbf{cyclist} : \langle e^a, t^a\rangle]\!]^{\mathcal{M}} =$ that function f from entities to truth values, defined so that, for all x, if $f(x) = 1$, then x is a cyclist, else $f(x) = 0$

This seems to be exactly what we want for **cyclist** when it is of type $\langle e^a, t^c\rangle$ as well. Only that assumption will account for the fact that one says Lance is a cyclist with both *Lance, a cyclist, is training* and *Lance is a cyclist.*

Thus, the factual support for the meaning language approach seems strong. Should we be suspicious on conceptual grounds? I argue that we should not. Although one hears continued talk of direct interpretation as a desirable framework, the talk does not match the practice of linguists. Arguably, every present-day framework that has been applied to a sufficiently wide range of data of both a syntactic and a semantic nature has ended up with crucial appeals to a meaning language in one form or another.

Some frameworks are open about this aspect of their design. Discourse Representation Theory (DRT) is explicitly a theory about semantic representations; DRT's boxes are the primary explanatory apparatus, and theorists in that tradition have worked hard to motivate this level of grammatical structure, often using variable binding and presupposition accommodation for this purpose (two examples are Kamp 1981 and van der Sandt 1992). The same is true of Glue semantics, where a resource-sensitive (linear) proof theory intervenes between the syntax (usually LFG functional-structures) and the semantics proper (lambda terms or the models). Here again, Glue theorists argue that it is a virtue; a recent argument involving ellipsis is Asudeh and Crouch 2002.

Other theories are less open about their use of a meaning language. For instance, semantic theories that seek to integrate the insights of transformational grammar (the Minimalist Program) interpret Logical Forms (LFs), which are sometimes claimed to be syntactic objects. But not since the late 1980s have these researchers taken seriously the idea that all covert movement operations must obey all known syntactic constraints. The consensus is that constraints on covert movement are vastly different from those on overt movement; examples of this position include Sauerland and Heck (2002) and Percus and Sauerland (2002). But this is just to say that the operations held to deliver the interpreted structure are not syntactic but semantic. LFs are in fact a meaning language. This perspective on them is greatly supported by the decision to include in LFs items that have the syntactic shape of lambda operators (Heim 1992), quantifiers (Heim 1982:§2, Chierchia 1998:366, Fox 2000), and semantic binding indices (Heim and Kratzer 1998). The idea is clear: we should be able to determine the look of the lambda term by reading off the leaves of the LF, using the structure to determine scopal relationships. The classic example of this is the calculus for arriving at LFs in Heim 1982:§2, in which LFs have the structure and the terminal nodes of first-order predicate-logic formulae. Since that time, LFs have come to resemble syntactic objects less and less.

The most forceful arguments for direct interpretation are found in the work of Jacobson (1999, 2000), which collect much of her earlier work on variable-free semantics. This theory has indeed done without necessary appeals to a meaning language. Instead, the syntactic part of the theory is responsible for the kind of combinatoric regulation that I have assigned to the type theory. One sees this mostly clearly in the role that superscripted syntactic categories play in Jacobson's work.

Jacobson (1999:129) defines a rule that takes any category A to a category A^{B}, where B is also a category. Thus, since S and NP are categories, S^{NP} is a category. In fact, this is the category of a sentence containing a deictic pronoun; the semantics of S^{NP} is that of a one-place function on individuals. The effect is, as Jacobson writes, to "mark the syntactic type of an expression so as to record its semantic type" (p. 129). One can fairly characterize these categories as *semantically informed syntactic categories.* In Shan (2001), they are relocated to the meaning language, as type distinctions.

Their primary purpose in both variants is to prevent unwanted combinatoric schemes. Jacobson (1999:128, 131–132) describes one of them: without the superscripting conventions, the VP *loves him* would be of category $\mathrm{S}/_{\mathrm{L}}\mathrm{NP}$, but have the semantics of a two-place relation. Assuming that *Mary* can be of category NP with the semantics of a referring expression, we could have derivations like the following:

(2.55)

$$\left\langle \begin{array}{c} \textit{mary loves him} \\ \mathrm{S} \\ \lambda y.\,\mathbf{love}(\mathbf{mary})(y) \end{array} \right\rangle$$

$$\left\langle \begin{array}{c} \textit{mary} \\ \mathrm{NP}\ ;\ \mathbf{mary} \end{array} \right\rangle \quad \left\langle \begin{array}{c} \textit{loves him} \\ \mathrm{S}/_{\mathrm{L}}\mathrm{NP} \\ \lambda x \lambda y.\,\mathbf{love}(x)(y) \end{array} \right\rangle$$

Though *Mary* is the subject in this sentence, its denotation ends up as the lovee (first argument to **love**). Moreover, the category of the resulting expression is S, but it has a functional (set-denoting) semantics, introducing a suspicious mismatch between the syntax and the semantics. The superscripting conventions correct this: *loves him* is actually of category $(\mathrm{S}/_{\mathrm{L}}\mathrm{NP})^{\mathrm{NP}}$. Thus, unlifted *Mary* is not an argument to this expression; *Mary* must type-raise and then undergo the g rule, which has the semantics of function composition and the syntactic effect of introducing a superscript onto the outermost argument slot. The details are not of significance here. What is important is the effect of the *syntactic* categories to regulate unwanted *semantic* composition schemes. With *Mary* properly shifted by the lift operation (to yield a generalized quantifier), and then shifted again by g (so that it has superscript marking), *Mary* combines with *loves him* in the desired fashion:

(2.56)

$$\left\langle \begin{array}{c} \textit{mary loves him} \\ \mathrm{S}^{\mathrm{NP}} \\ \lambda x.\, \mathbf{love}(x)(\mathbf{mary}) \end{array} \right\rangle$$

$$\left\langle \begin{array}{c} \textit{mary} \\ \mathrm{S}^{\mathrm{NP}}/_{\mathrm{R}}(\mathrm{S}/_{\mathrm{L}}\mathrm{NP})^{\mathrm{NP}} \\ \lambda F \lambda x.\, (\lambda f.\, f(\mathbf{mary}))(F(x)) \end{array} \right\rangle \quad \left\langle \begin{array}{c} \textit{loves him} \\ (\mathrm{S}/_{\mathrm{L}}\mathrm{NP})^{\mathrm{NP}} \\ \lambda y \lambda z.\, \mathbf{love}(y)(z) \end{array} \right\rangle$$

$$\mathbf{g} \quad \left\langle \begin{array}{c} \textit{mary} \\ \mathrm{S}/_{\mathrm{R}}(\mathrm{S}/_{\mathrm{L}}\mathrm{NP}) \\ \lambda f.\, f(\mathbf{mary}) \end{array} \right\rangle$$

$$\mathbf{lift} \quad \left\langle \begin{array}{c} \textit{mary} \\ \mathrm{NP} \\ \mathbf{mary} \end{array} \right\rangle$$

The example is typical; the motivation for the superscripting convention is largely its ability to regulate semantic composition in useful ways. I do not mean to imply that researchers in the variable-free framework should place these distinctions in the meaning language instead. Rather, I wish to point out only that the ideas presented in the present work are compatible with direct interpretation as Jacobson conceives of it. In fact, it is worth pointing out that a consequence of the meaning language distinctions made here is that we can interpret surface structures without need of new syntactic apparatus or movement operations. Since the major impetus of Jacobson's variable-free program seems to be that "surface structures directly receive a model-theoretic interpretation without being mapped into another level (i.e., LF)" (Jacobson 1999:117), the theory of CIs fits well into this overall research program, with a slight caveat: to fully comply with Jacobson's tenets, we move the meaning language distinctions into the (categorial) syntax, as semantically informed syntactic categories.

2.8 Discourse structures

At this point, I shift attention from translations to model-theoretic interpretations. An important feature of the definition of CIs is that they are always speaker-oriented. This is straightforwardly captured as part of how the interpretation function works, as specified with **parsetree interpretation** (2.50). Since we are working with a system of indirect translation, a translation function, $\rightsquigarrow$, takes natural language expressions to $\mathcal{L}_{\mathrm{CI}}$. The interpretation brackets take terms of $\mathcal{L}_{\mathrm{CI}}$ to model-theoretic entities.

The goal is to interpret discourses, so we should let the intuitive notion of what a discourse is be our guide in constructing the models. A discourse consists of a set of discourse participants. Each might view the world quite differently from the others, but they share a language. When a discourse participant a makes an utterance, the other discourse participants learn something about the way that a's model of the world looks. The next few subsections are devoted to developing structures that contain these elements. I call them *discourse structures.*

2.8.1 The discourse layer

As a first step towards discourse structures, let's look at (2.57a) and its somewhat stilted counterpart (2.57b).

(2.57)
a. Jed said, "Ed fled".
b. Jed bears the utterance relation to the sentence *Ed fled.*
c. When Jed said "Ed Fled", I was out by the shed.
d. There is a past time t such that Jed uttered the sentence "Ed Fled" at t and I was in the shed at t.

The examples in (2.57c)–(2.57d) are more elaborate illustrations of the same basic fact: agents can enter into relations with sentences, and utterances are things that we can locate in space–time and assign properties to. Needed, then, is the ability talk about both sentences and individuals as entities in the semantic model. So let's give ourselves that ability. The following clauses provide most of the required tools.

(2.58)

a. $\mathcal{L}_U$ is a higher-order lambda calculus with types e, t, d, and q, where $d \cup q = u$. (This logic is defined in appendix A.)

b. D_e, the domain of type e, is a set of entities.

c. $A = \{a_1, a_2, \ldots\}$ is a set of discourse participants; $A \subseteq D_e$.

d. $D_u = \{\mathcal{S}_1, \mathcal{S}_2, \ldots\}$ is a set of sentences, the domain of u. Each $\mathcal{S}$ is a pair $(\mathcal{T}^s, \mathcal{T}^m)$, in which $\mathcal{T}^s$ is a syntactic structure and $\mathcal{T}^m$ is its associated semantic parsetree (as defined in (2.36)). D_u contains a subset $D_d = \{\mathcal{D}_1, \mathcal{D}_n, \ldots\}$ of declaratives (the domain of d) and a subset $D_q = \{\mathcal{Q}_1, \mathcal{Q}_n, \ldots\}$ of interrogatives (the domain of q). $D_q \cap D_d = \emptyset$.

e. D_t, the domain of t, is $\{1, 0\}$, the set of truth values.

f. $V_{\mathcal{D}}$ is a valuation function, taking constants of $\mathcal{L}_U$ to functions formed from objects in $D_e \cup D_u \cup \{0, 1\}$, constrained so that if α is of type σ, then $V_{\mathcal{D}}(\alpha) \in D_\sigma$.

Let's say that, in $\mathcal{L}_U$, terms are always given inside corner brackets. Thus, $\ulcorner$**frankly**$\urcorner$ and $\ulcorner$**utter**$\urcorner$ are terms of $\mathcal{L}_U$. Additionally, I adopt the convention that the $\mathcal{L}_U$ constant associated with a sentence $\mathcal{S} \in D_u$ is always the surface realization of $\mathcal{S}$ with corner brackets around it. For example:

(2.59) $V_{\mathcal{D}}(\ulcorner\text{Ed fled}\urcorner) =$

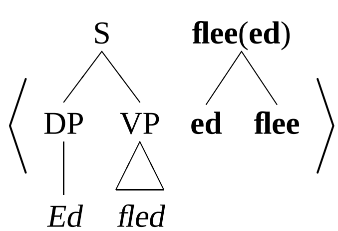

With these basic tools, we can now provide a semantics for expressions like (2.57a). Let $\ulcorner$**utter**$\urcorner$ be a term of $\mathcal{L}_U$ of type $\langle d, \langle e, t\rangle\rangle$, where d is the type of expressions denoting in D_d. That is, $\ulcorner$**utter**$\urcorner$ names a functions from declarative sentences to functions from entities to truth values. We then have the following analysis of (2.57a).

(2.60) *Jed said, "Ed fled"* $\rightsquigarrow$ $\ulcorner$**utter**$\urcorner$($\ulcorner$Ed fled$\urcorner$)($\ulcorner$**jed**$\urcorner$) : t

2.8.2 The lower layer

The logic $\mathcal{L}_U$ is inadequate in one important sense: it treats sentences as though they were atomic objects. Both $\ulcorner$Ed fled$\urcorner$ and $\ulcorner$Ali ran the marathon$\urcorner$ are simply constants, denoting entities in D_u. We need some way to interpret members of D_u. That is, we need some way of doing the usual things we do in semantics.

To achieve this added dimension to the logic and its models, I appeal to layering techniques. The new pieces are given in (2.61)–(2.62).

(2.61) $\mathfrak{M} = \{\mathcal{M}_1, \mathcal{M}_2, \ldots\}$ is a set of intensional models for the logic $\mathcal{L}_{\text{CI}}$. Each $\mathcal{M}_i \in \mathfrak{M}$ is a pair $(\mathbf{D}, V_i)$, where

a. $\mathbf{D}$ is a set of domains, common to all models in $\mathfrak{M}$ and defined as follows:

   i. The domain of e^a and e^c is D_e, a set of entities.
   ii. The domain of s^a and s^c is D_s, a set of entities called worlds, disjoint from D_e.
   iii. The domain of t^a and t^c is $D_t = \{0, 1\}$, the set of truth values.
   iv. The domain of a functional type $\langle\sigma, \tau\rangle$ is $\{f \mid f : D_\sigma \mapsto D_\tau\}$.
   v. The domain of a product type $\langle\sigma \times \tau\rangle$ is $D_{\langle\sigma\times\tau\rangle} = D_\sigma \times D_\tau$.

b. V_i is a valuation taking formulae of $\mathcal{L}_{\text{CI}}$ to the model, constrained so that if $\alpha \in ME_\sigma$, then $V_i(\alpha) \in D_\sigma$.

(2.62) $\hbar$ is a function that takes each $a_i \in A$ to the model $\mathcal{M}_i \in \mathfrak{M}$, where $\mathcal{M}_i$ can be viewed as the world-view of a_i

The set $\mathfrak{M}$ provides us with world-views for our discourse participants. The function $\hbar$ associates each discourse participant with his or her own intensional model. The model $\mathcal{M}_i$ represents the world-view of the discourse participant a_i. These intensional models are total (fully specified), in the sense that their valuations are total functions. This is a fiction — speakers do not in fact know, or pretend to know, everything — but it is a convenient one. We could restate the definitions in terms of partial structures, but the complications would be considerable.

2.8.3 Interpretation

I've essentially defined two logics and two classes of models. They are brought together by the interpretation function, which is sensitive to the sort of object it is applied to. In essence, we can feed it either expressions of $\mathcal{L}_U$ or parsetrees in the set specified in (2.36). First, I bring together the above ideas in a definition of *discourse structure*:

(2.63) A *discourse structure* is a tuple $\mathcal{D} = (A, \mathbf{D}, D_u, \mathfrak{M}, \hbar, V_{\mathcal{D}})$, where

a. $A = \{a_1, a_2, \ldots\}$ is a set of discourse participants.

b. $D_u = \{\mathcal{S}_1, \mathcal{S}_2, \ldots\}$ is a set of sentences, the domain of u. Each $\mathcal{S}$ is a pair $(\mathcal{T}^s, \mathcal{T}^m)$, in which $\mathcal{T}^s$ is a syntactic structure and $\mathcal{T}^m$ is its associated semantic parsetree (as defined in (2.36)). D_u contains a subset $D_d = \{\mathcal{D}_1, \mathcal{D}_n, \ldots\}$ of declaratives (the domain of d) and a subset $D_q = \{\mathcal{Q}_1, \mathcal{Q}_n, \ldots\}$ of interrogatives (the domain of q). $D_q \cap D_d = \emptyset$.

c. $\mathbf{D}$ is a set of domains, as defined in (2.61); $A \subseteq D_e$.

d. $\mathfrak{M} = \{\mathcal{M}_1, \mathcal{M}_2, \ldots\}$ is a set of intensional models, as defined in (2.61). All $\mathcal{M}_i \in \mathfrak{M}$ have $\mathbf{D}$ as their set of domains.

e. $\hbar$ is a function that takes each $a_i \in A$ to the model $\mathcal{M}_i \in \mathfrak{M}$, where $\mathcal{M}_i$ can be viewed as the world-view of a_i.

f. $V_{\mathcal{D}}$ is a valuation function, taking constants of $\mathcal{L}_U$ to functions formed from objects in $D_e \cup D_u \cup \{0, 1\}$, constrained so that if α is of type σ, then $V_{\mathcal{D}}(\alpha) \in D_\sigma$.

The definition assumes that all discourse participants "speak" (form meanings from) the same meaning language, $\mathcal{L}_{\mathrm{CI}}$. The fact that each intensional model in a discourse structure has the same domains $\mathbf{D}$ ensures that everyone is talking about the same stuff. These too are simplifying measures, permitting easier definition of notions like *common ground* should they prove useful. (The common ground is that valuation function V_c such that $V_c(\varphi)$ is defined only if, for all discourse participants a_i, a_j, we have $V_i(\varphi) = V_j(\varphi)$.)

The interpretation function $[\![\cdot]\!]^{\mathcal{D},s,a}$ for a discourse structure $\mathcal{D}$ is relativized to a speaker s and an addressee a, both members of A. It is defined as in (2.64).

(2.64) a. $[\![\varphi]\!]^{\mathcal{D},s,a} = V_{\mathcal{D}}(\varphi)$ if φ is a formula of $\mathcal{L}_U$.

b. $[\![\mathcal{S}]\!]^{\mathcal{D},s,a}$ = the value of $\mathcal{S}$ determined by (2.50) if $\mathcal{S}$ is a parsetree for $\mathcal{L}_{\mathrm{CI}}$.

We are now positioned to return to the basic case (2.57a), to find out just what it is that Jed asserted when he entered into this particular relation with the sentence *Ed fled.* First, we apply the interpretation function to the translation in (2.60).

(2.65) $$[\![\ulcorner\textbf{utter}\urcorner(\ulcorner\text{Ed fled}\urcorner)(\ulcorner\textbf{jed}\urcorner)]\!]^{\mathcal{D},s,a} = 1 \text{ iff}$$
$$[\![\ulcorner\textbf{utter}\urcorner]\!]^{\mathcal{D},s,a}\left(\left\langle \left[_{\text{S}}\ [_{\text{DP}}\ \textit{Ed}]\ [_{\text{VP}}\ \textit{fled}]\right]\quad \left[_{\textbf{flee(ed)}}\ \textbf{ed}\ \ \textbf{flee}\right] \right\rangle\right)([\![\ulcorner\textbf{jed}\urcorner]\!]^{\mathcal{D},s,a}) = 1$$

Here is what Jed said when he said *Ed fled*:

(2.66) $$\left[\!\!\left[\left\langle \left[_{\text{S}}\ [_{\text{DP}}\ \textit{Ed}]\ [_{\text{VP}}\ \textit{fled}]\right]\quad \left[_{\textbf{flee(ed)}}\ \textbf{ed}\ \ \textbf{flee}\right] \right\rangle \right]\!\!\right]^{\mathcal{D},\text{Jed},a} = [\![\textbf{flee}(\textbf{ed})]\!]^{\mathcal{M}_{\text{Jed}},g} = 1 \text{ iff Ed fled in } \mathcal{M}_{\text{Jed}}$$

That is, Jed informed the other members of A that, in his model $\mathcal{M}_{\text{Jed}}$, it is true that Ed fled. This shift in the speaker parameter (from s to Jed) is, I assume, engendered by the presence of the quotation marks.

The upper layer of these discourse structures permits us to talk about, and place conditions on, discourses. This ability proves essential for the discussion of utterance modifiers like *frankly* in chapter 3, section 3.7.3, and the discussion of performative honorifics in chapter 4, section 4.6.2. The nature of the interpretation function also provides a suitable formalization of speaker-orientation. If I say, "Ed fled", I stand in the utterance relation to the sentence *Ed fled.* When we interpret the content of this sentence, we do so in the intensional model $\mathcal{M}_{\text{chris}}$. This provides a precise characterization of the speaker-orientation of main clause assertions. By the rule of **parse-**

tree interpretation, (2.50), all CI propositions are also interpreted in $\mathcal{M}_{\text{chris}}$, thereby achieving speaker-orientation for them as well.

2.9 The heritage function

A prominent feature of Karttunen and Peters' (1979) semantics that I do not adopt here is their *heritage function*. The purpose of this function, which is essentially a third dimension to lexical entries, is to regulate the way that (what they define as) CIs are inherited in complex sentences. The view of CIs I advocate removes all need for a heritage function. This section explores why this is so.

As I noted in chapter 1.2, the heritage function is really a presupposition projection function under a new name, and subsequent authors have treated it as a variant of that general mode of theorizing. This is just one of the respects in which Karttunen and Peters make a terminological adjustment. Throughout this section, I use the locution 'K&P-CI' to refer to their definition of conventional implicature. For the most part, 'K&P-CI' is coextensive with 'presupposition'.

The factual justification that Karttunen and Peters (1979) offer for a heritage function derives entirely from presupposition triggers. For instance, they claim that the heritage function for factive verbs like *forget* (with tensed complement) and *realize* contrasts with the heritage function of propositional attitude verbs and verbs of saying (p. 20ff). They provide examples like those in (2.67), in which the cleft in the complement triggers the K&P-CI proposition that Mary's phone was tapped by someone.

(2.67) a. It wasn't Bill who tapped Mary's phone.

b. John forgot that it wasn't Bill who tapped Mary's phone.

c. John believed that it wasn't Bill who tapped Mary's phone.

They observe that (2.67a) and (2.67b) share the K&P-CI that Mary's phone was tapped, whereas (2.67c) "conventionally implicates only that John believed someone tapped Mary's phone; the speaker of the sentence does not commit himself to the belief's being correct, but only to John's having had it" (p. 20).

Thus, in Karttunen and Peters' view, the verbs *forget* and *believe* have different heritage functions. The heritage function of **forget** is an identity function for both at-issue and CI content, whereas **believe** relativizes the at-issue and CI content of its first argument to the beliefs of its second argument. Similarly, they adopt, in their section 6, a set of heritage principles for compound sentences that essentially reencodes the presupposition inheritance principles of Karttunen 1973.

The brief quotation just above indicates why the heritage function is not needed. Karttunen and Peters say that the CIs of the complement to *believe* are relativized to the subject of the belief predication. On the view taken here, CIs are invariably speaker-oriented; a heritage function for conventional implicatures would always be an identity function. Consider the following variations on the above examples:

(2.68) a. Bill is an advocate of individual privacy laws.

b. John forgot that Bill, an advocate of individual privacy laws, tapped Mary's phone.

c. John believed that Bill, an advocate of individual privacy laws, tapped Mary's phone.

Even in the case of *believe*, the CI represented by the content of (2.68a) is inherited. If we adopted Karttunen and Peters' heritage function, we would lose this result — the meaning of *believe* would filter off the content of the nominal appositive, reducing it to a mere belief of Bill's.

One could of course adopt something like a heritage function, sticking close to the original two-dimensional semantics. But it would invariably be an identity function. Moreover, we would require some way to accumulate CIs as the composition proceeded. The result would be that the denotation of the root node of a tree would contain a tuple of CIs — all and only the CIs determined by the lexical content of the tree in question. We can call this a 'widest-scope strategy', because it seeks to ensure that every CI always has the widest scope possible — out of all syntactic islands, intensional contexts, etc.

But the relevant concept is not widest scope, but rather *scopelessness*. It is for this reason that I rejected, in section 2.6.7, this method of handling accumulated CIs. The intuition behind getting them all to the top of the tree is just that they must be prevented from interacting with any operators that have syntactic scope over the items that generated them. The surest way to do this, I argue, is essentially to remove them from the at-issue semantics entirely, and as soon as they are propositional. The rule of **parsetree interpretation** defined in (2.50) achieves this.

2.10 The 'binding' problem (or virtue)

Beaver (1997, 2001) identifies what he calls the 'binding problem' with Karttunen and Peters' (1979) two-dimensional semantics: there is no way to bind a variable in both dimensions of meaning with the same quantifier (see also Cooper 1983:151–152; Heim

1983; van der Sandt 1992:338–339; Krahmer 1998:12, 21, 118ff). Thus, Karttunen and Peters (1979:54) recognize that their semantics for (2.69a) is (2.69b), in which there are two distinct existential quantifiers:

(2.69) a. Someone managed to succeed George V.

b. $$\begin{bmatrix} \exists x[\textbf{succeed}(\textbf{george})(x)] \\ \bullet \\ \exists y[\textbf{try-hard-to}(\textbf{succeed}(\textbf{george})(y))(y)] \end{bmatrix}$$

This is an incorrect result. Suppose we interpret this relative to a model $\mathcal{M}$ in which the successor of George V had an easy time assuming the throne, but a deranged peasant struggled to obtain the throne through a grassroots campaign proclaiming himself the true son of George V. We want (2.69a) to be false in $\mathcal{M}$, but (2.69b), its hypothesized translation, is true in $\mathcal{M}$.

A reasonable reaction to this failure is to deny that *manage* and other implicative verbs contribute in terms of CIs. The 'binding problem' counts as a problem only relative to a particular set of facts; it is not a *logical* problem. In the context of the current work, this view of the limitation is easily supported: implicative verbs behave like presupposition triggers when embedded below presupposition plugs (Karttunen 1971; Karttunen and Peters 1979:20), whereas CIs do not do this, at least as conceived of here. Implicatives seem also to differ from CIs in terms of their discourse requirements; though Karttunen and Peters (1979) claim that the non-at-issue content of implicative verbs is neither backgrounded nor required to be true for definedness of the whole, most speakers disagree with these judgments.

However, in the present setting, we need not lean to hard on the nature of the judgments to provide an analysis of these cases that squares with the independence of the at-issue and CI dimensions. We can either define *manage* as a presupposition trigger (Krahmer 1998:§5), or we can make use of our product types, which are formed only from at-issue meanings. We could, for instance, obtain an adequate description of (2.69a) as follows:

(2.70) a. *someone* $\rightsquigarrow$
$\lambda\langle f, g\rangle.\, \exists x[\langle f(x), g(x)\rangle] : \langle \overrightarrow{\langle e^a, t^a\rangle}_2, \overrightarrow{t^a}_2\rangle$

b. *manage* $\rightsquigarrow$
$$\lambda f. \left\langle \begin{array}{c} \lambda y.\, f(y), \\ \lambda z.\, \textbf{try-hard-to}(f(z))(z) \end{array} \right\rangle : \langle\langle e^a, t^a\rangle, \overrightarrow{\langle e^a, t^a\rangle}_2\rangle$$

c.
$$\exists x \left[\left\langle \begin{array}{c} \textbf{succeed}(\textbf{george})(x), \\ \textbf{try-hard-to}(\textbf{succeed}(\textbf{george}))(\textbf{george}) \end{array} \right\rangle \right] : \overrightarrow{t^a}_2$$

$\textbf{someone} : \langle \overrightarrow{\langle e^a, t^a\rangle}_2, \overrightarrow{t^a}_2\rangle$

$\textbf{manage}(\textbf{succeed}(\textbf{george})) : \overrightarrow{\langle e^a, t^a\rangle}_2$

$\textbf{manage}$ $\langle\langle e^a, t^a\rangle, \overrightarrow{\langle e^a, t^a\rangle}_2\rangle$

$\textbf{succeed}(\textbf{george}) : \langle e^a, t^a\rangle$

For the range of facts under discussion here, the 'binding problem', rather than being a detriment to the system, is actually a virtue of it. As observed in Potts 2002b, it is impossible to bind a variable into an *As*-clause:

(2.71) *No reporter$_1$ believes that, as he$_1$ wrote, Ames is a spy.

There is no interpretation of this in which *no reporter* functions as a semantic binder for the pronoun *he* in the *As*-clause. But it is easy to see why this is so; consider the parsetree for this example, given in (2.72), using the current combinatorics and style of lexical denotations.

(2.72)

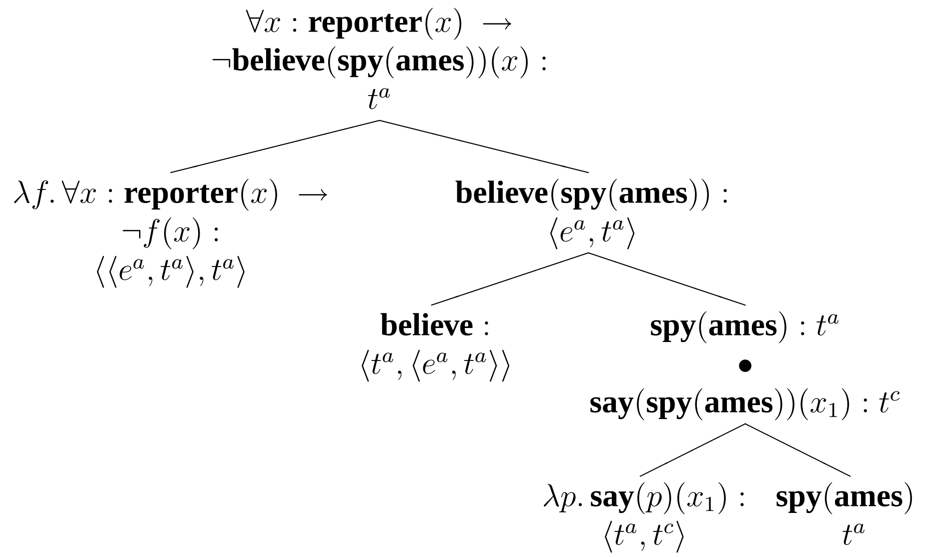

Importantly, the variable x_1 in the *As*-clause meaning $\lambda p.\, \textbf{say}(p)(x_1) : \langle t^a, t^c\rangle$ is not bound by the universal quantifier; it is therefore interpreted as a free pronoun. This is a direct result of the fact that the two variables are in different dimensions. (Any occurrence of x_1 in the at-issue meaning would get bound, of course.) Thus, the impossibility of binding into this supplement is a direct result of the logic. One might think that it traced back to the lexical meaning of the quantifiers, which are defined to take at-issue values into same. This is true, but the limitation runs deeper: a quantifier that could take a CI as its argument to return an at-issue meaning would have to be of extensional type $\langle\langle e^a, t^c\rangle, t^a\rangle$. But this is not a member of the set of types in (2.18).

The observation is not unique to *As*-clauses. The other supplements under study here fail to allow variable-binding from outside them. I illustrate in (2.73) and (2.74), using the negative universal quantifier *no* to ward off potential E-type readings of the variables, readings which are usually not available for downward entailing operators (Fox 2000:56ff).

(2.73) *No reporter$_1$ believes that Ames is a spy, which he$_1$ wrote in his column.

(2.74) *No reporter$_1$ believes that Ames, often the subject of his$_1$ columns, is a spy.

The current logic derives this failure from unexceptional facts about quantifier scope. I undertake more detailed discussions of this point in chapters 3 and 4.

2.11 One-dimensional translations

The generalization in (2.2), repeated here, isolates a property of the set of CI-based lexical items discussed in this work:

(2.75) No lexical item contributes both an at-issue and a CI meaning.

This seems at odds with past discussions of CIs. In particular, it is quite different from the fragment of Karttunen and Peters (1979), in which all expressions come with both dimensions of meaning. However, as noted in the introduction, their proposal does not directly conflict with (2.75). There are three cases to consider. All are in fact in harmony with (2.75).

In the first case, one of the dimensions of meaning is an identity function. It is somewhat problematic for me to use Karttunen and Peters' examples directly, since they consist mainly of presupposition triggers or purely at-issue operators. But the kind of meaning that is of concern is represented by the alternative meaning for an expressive attributive adjective in (2.76) (which, like the lexical entry in chapter 4, uses the function called *kind*, symbolized $^{\cap}$, which takes properties to their individual-level correlates).

(2.76) $\quad$ *damn* $\rightsquigarrow \begin{bmatrix} \lambda f.\, f \ : \ \langle\langle e^a, t^a\rangle, \langle e^a, t^a\rangle\rangle \\ \bullet \\ \lambda f.\, \textbf{bad}(^\cap f) \ : \ \langle\langle e^a, t^a\rangle, t^c\rangle \end{bmatrix}$

This appears to have two dimensions of meaning. But the at-issue dimension is an identity function; it has no effect whatsoever on the information content of the whole. Its presence obscures the generalization in (2.75). This is one of the reasons for placing the effects of the identity function in the general rules for combining meanings. Rather than having structures like (2.77), we have the more transparent (2.78), which employs the rule **CI application** defined in (2.42) above.

(2.77)

$$
\begin{array}{c}
\textbf{republican} : \langle e^a, t^a\rangle \\
\bullet \\
\textbf{bad}(^\cap\textbf{republican}) : t^c
\end{array}
$$

$$
\begin{array}{cc}
\begin{array}{c}
\lambda f.\, f : \langle\langle e^a, t^a\rangle, \langle e^a, t^a\rangle\rangle \\
\bullet \\
\lambda f.\, \textbf{bad}(^\cap f) : \langle\langle e^a, t^a\rangle, t^c\rangle
\end{array}
&
\textbf{republican} : \langle e^a, t^a\rangle
\end{array}
$$

(2.78)

$$
\begin{array}{c}
\textbf{republican} : \langle e^a, t^a\rangle \\
\bullet \\
\textbf{bad}(^\cap\textbf{republican}) : t^c
\end{array}
$$

$$
\begin{array}{cc}
\lambda f.\, \textbf{bad}(^\cap f)) : \langle\langle e^a, t^a\rangle, t^c\rangle
&
\textbf{republican} : \langle e^a, t^a\rangle
\end{array}
$$

A second class of cases concerns content that is actually conversationally implicated (and hence deniable). Some examples of this that appear in Karttunen and Peters 1979 are offered by them merely as expository devices, but it is worth looking at them anyway. One example concerns the proper name *Bill*, which they assign the following semantics (p. 52):

(2.79) $\quad$ *Bill* $\rightsquigarrow \left[\begin{array}{c} \lambda f.\, f(\mathbf{bill}) : \langle\langle e^a, t^a\rangle, t^a\rangle \\ \bullet \\ \lambda f.\, \mathbf{male}(\mathbf{bill}) : \langle\langle e^a, t^a\rangle, t^c\rangle \end{array}\right]$

This meaning is suspicious first of all because the second meaning is not really functional at all. The abstraction is vacuous; the term is equivalent to $\mathbf{male}(\mathbf{bill}) : t^c$. But it suffers factual problems as well. By making maleness an *entailment* of the name *Bill*, Karttunen and Peters render sentences like (2.80) not just pragmatically surprising given our society's naming conventions, but in fact semantically contradictory.

(2.80) $\quad$ Bill is a girl.

Since this is not tautologically false, but rather merely unusual, the proper classification of maleness is as a conversational implicature.

The third and final apparent challenge to (2.75) involves presuppositions. I consider here their denotations for the universal quantifier (p. 49), which is given in (2.81).

(2.81) $\quad$ *every* $\rightsquigarrow \left[\begin{array}{ccc} \lambda f \lambda g.\, \forall x[f(x) \to g(x)] & : & \langle\langle e^a, t^a\rangle, \langle\langle e^a, t^a\rangle, t^a\rangle\rangle \\ & \bullet & \\ \lambda f \lambda g.\, \exists x[f(x) \wedge g(x)] & : & \langle\langle e^a, t^a\rangle, \langle\langle e^a, t^a\rangle, t^c\rangle\rangle \end{array}\right]$

In this case, assuming a classical universal quantifier, both dimensions contribute nontrivial information. But the consensus is that the CI meaning in (2.81) is actually a presupposition (or perhaps a conversational implicature). One reason to assume this is that it is a cancellable meaning. Suppose, for instance, that an airline implements the following regulation:

(2.82) $\quad$ Every airplane with pets on board must obtain a special permit from flight control.

Suppose that my job is to ensure that (2.82) is complied with every day. One day, no planes take off with pets on board. Am I entitled to say that (2.82) was complied with? Of course I am. The regulatory context conspires to cancel the usual existential presupposition. If we made that presupposition a CI, it would become an entailment. The result would be the incorrect prediction that I must answer that (2.82) was not complied with on any pet-free day.

In section 1.4.5, I addressed the semantics of *but*, another potential exception to (2.75), and showed that it is best treated as contributing entirely in terms of at-issue content, a view that is in line with Bach (1999). In sum, we're justified in concluding that Karttunen and Peters 1979 does not contain evidence that is contrary to the generalization in (2.75).

Having argued that there are not serious challengers to (2.75), I should ask whether it is an isolated stipulation or whether it is a consequence of anything about the description logic. The statement that we require strikes me as unsurprising. It is simply this: lexical meanings are terms of $\mathcal{L}_{\mathrm{CI}}$. The definition of the set of well-formed expressions in (2.20) (given in complete detail in appendix A) contains no terms of the form

$$\begin{array}{c} \alpha : \sigma^a \\ \bullet \\ \beta : \tau^c \end{array}$$

These can arise only in parsetrees, as the result of combining more basic expressions.

2.12 A note on resource sensitivity

Klein and Sag (1985) venture a simple but ultimately quite significant hypothesis about natural language semantics. They write, (2.83) "Translation rules in Montague semantics have the property that the translation of each component of a complex expression occurs exactly once in a translation of the whole. [...] This is to say, we do not want the set S [of lexical items —C. P.] mentioned above to contain *all* meaningful expressions of IL which can be built up from elements of S, but only those which use each element of S exactly once."

(Klein and Sag 1985:172)

The worry driving (2.83) is easy to illustrate. Suppose we have the set of lexical items $\{\textit{Ed}, \textit{catch}, \textit{fish}\}$. Each translates as the name of a function: $\{\mathbf{ed}, \mathbf{catch}, \mathbf{fish}\}$. Absent the injunction in (2.83), nothing prevents a derivation in which **ed** is used twice and *fish* not at all: from the sentence *Ed caught fish*, one could derive a lambda term expressing the proposition that Ed caught Ed.

Klein and Sag (1985) proceed to define a notion of 'bounded closure' that ensures this 'all and only' quality. Similar concepts appear throughout linguistics. For instance, in the syntactic framework of the Minimalist Program, the numeration (a multiset of lexical items) must be exhausted in the course of a derivation — *all* lexical items must be used. The inclusiveness principle (Chomsky 1995:228) seeks, in part, to limit derivations to *only* items in the initial numeration. While Chomsky (1995:228) says that 'the inclusiveness condition is not fully met", the intuition remains a guiding force in that tradition. What's more, Brody (1995) observes that, in the copy theory of movement, a condition must ensure that exactly one member in a copy chain is interpreted. The overarching generalization is that 'bounded closure' has a role to play in linguistic systems. Asudeh (2003) offers a fuller discussion of these examples and others.

Researchers in Glue semantics regard this situation as a signal that the logic underlying linguistic theories is not a classical one. In classical systems, a single instance of

p as a premise can be used multiple times: p implies $(p \wedge p)$, and so forth (right weakening). Conversely, multiple instances of p imply a single instance of p — for example, $(p \wedge p)$ classically entails p (left weakening). Denying both inferences results in *linear logic* (or a sublogic thereof), which has the 'once and only once' character described in (2.83). In linear logic, and in turn in Glue semantics (a multiplicative fragment of linear logic), meanings are literally consumed.

The composition rule in (2.42) seems to challenge the resource-sensitivity premise. I repeat the rule here:

(2.84) **CI application**

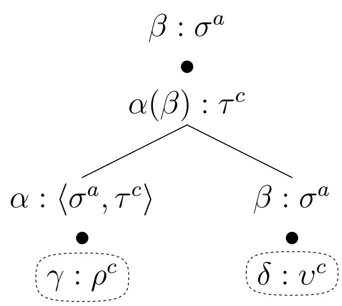

The meaning β is used (consumed) twice. Such reuse is essential to the logic for CIs. We can see this even more clearly if we view (2.84) as a proof rule. The procedure is relatively simply. First, remove the optional terms, as they are not relevant to the nature of the rule. Second, turn the tree admissibility over and replace the dominance lines with a horizontal one. Finally, using an informal version of the Curry–Howard bijection, replace all basic types with propositional letters $p, q, \ldots$, and replace all functional types with implications like $p \to q$. We arrive at the natural deduction proof rule in (2.85).

(2.85) $$\frac{p \qquad p \to q}{p \quad q}$$

I conclude from this not that we need to rework $\mathcal{L}_{\mathrm{CI}}$, but rather that the resource-sensitivity premise of these logics is too strong. It seems to me that the limited sort of reuse represented by (2.84) captures the essence of the role CI expressions play in natural language.

But this conclusion is not inevitable. Though (2.84) indeed violates the tenets of multiplicative linear logic and categorial grammar, there are alternatives that are in keeping with these logics. Ash Asudeh (p.c., 1/03) offered a solution that depends on having a premise that performs the needed duplication, so that we have no reuse at all. The work of **CI application** is done by individual lexical items.

We should look closely at Asudeh's proposal. $\mathcal{L}_{\mathrm{CI}}$ shares important properties with the Glue semantics. In particular, $\mathcal{L}_{\mathrm{CI}}$ depends upon the Glue semantics premise that proofs (here, semantic parsetreees) have a nontrivial role to play in semantic theory (Katz and Katz 1977; Asudeh 2003). The following is a modified version of Asudeh's proposal. (I have made a few changes; Asudeh might not endorse this version of his ideas.)

First, we import the type-theoretic distinctions of $\mathcal{L}_{\mathrm{CI}}$ into the Glue setting. To do this, I assume that our propositions (*meaning constructors*, in Glue parlance) are marked as either at-issue or CI. Thus, we have p^a and p^c, leaving open that these might denote the same objects. We can form complex expressions from these basic meaning constructors using the resource sensitive implication operator $\multimap$ (roughly equivalent to the basic type constructor of $\mathcal{L}_{\mathrm{CI}}$, $\langle\cdot,\cdot\rangle$), as well as the resource sensitive coordinator $\otimes$. With these distinctions in place, we can have analyses like the following:

(2.86) a. *damn* $\rightsquigarrow f^a \multimap (f^a \otimes p^c)$

b. *Republicans* $\rightsquigarrow f^a$

c. *damn Republicans* $\rightsquigarrow$ $$\dfrac{\dfrac{f^a \qquad f^a \multimap (f^a \otimes p^c)}{f^a \otimes p^c}}{f^a \quad p^c}$$

The key element is the meaning constructor for *damn*, which takes a f^a and returns a pair of meaning constructors, which we can then split apart and treat independently. The result is a proof that terminates with a pair of meaning constructors (only this kind of multidimensionality can satisfies the dimensional independence properties uncovered above). But this seems like an acceptable modification to the usual way of theorizing in Glue. The important thing is that we have not reused any premises; the **CI application** rule is replaced by individual lexical items with the property that they take a given meaning and return that same meaning coordinated with something else.

We can move fairly easily between $\mathcal{L}_{\mathrm{CI}}$ and the Glue semantics perspective. The differences take the form of metagrammatical considerations about how one wishes to handle the complexity of CI expressions. In Glue semantics, the logical novelties of CIs are located in the lexicon; we understand the way CI expressions work, not by studying the nature of the inference rules, but rather by studying the lexical items with meaning constructors of the form $a^a \multimap (a^a \otimes p^c)$. In the $\mathcal{L}_{\mathrm{CI}}$ treatment, the special properties of CIs are readily apparent in the design of the logic. The type theory divides up the lexicon, and the limited reuse rule **CI application** brings to the fore the way that CI expressions borrow from the at-issue dimension without creating any kind of deficit there.

2.13 Chapter summary

I close this chapter by making explicit the connections between the description logic $\mathcal{L}_{\text{CI}}$ and the original definition of CIs. My review takes the form of a trip through the definition, with the numbering from chapter 1.

(1.10a) CIs are part of the conventional (lexical) meaning of words.

This is a consequence of the fact that terms of the form $\alpha : \sigma^c$ are the translations of certain lexical entries. The general combinatoric rules tell us how to manipulate them in structures.

The next clause is immediate from the multidimensional view of meaning:

(1.10b) CIs are commitments, and thus give rise to entailments.

Some nodes (all of them nonterminals) in the parsetrees for $\mathcal{L}_{\text{CI}}$ are decorated with two terms. The CI collection rule ensures that if α is of type $\langle s^a, t^c \rangle$ (extensionally, t^c) then α is interpreted as an entailment of the sentence.

The next clause is intimately related to the previous one:

(1.10c) These commitments are made by *the speaker of the utterance* “by virtue of the meaning of” the words he chooses.

When we interpret a formula of $\mathcal{L}_{\text{CI}}$ relative to a discourse structure, it is always relative to one of the discourse participants, identified as the speaker. Since CI propositions cannot combine with higher intensional operators, they never end up interpreted relative to any index but the one specified as the speaker parameter on the interpretation function. The effect is that, no matter how deeply embedded syntactically, CIs are always interpreted as though they were root-level assertions.

With the final clause, we return once again to the heart of this approach: the independence of the two dimensions of meaning:

(1.10d) CIs are logically and compositionally independent of what is “*said* (in the favored sense)”, i.e., independent of the at-issue entailments.

The rule of **CI application**, (2.42), permits a CI term to apply to an at-issue input. But the result is always a pair of terms, each interpreted independently by **parsetree interpretation**, (2.50). Moreover, the at-issue input is also passed on unmodified; if we were to snip off all CI terms from a parsetree, we would find its at-issue value unchanged.

Chapter 3

Supplements

3.1 Remarks

The general characteristics of *supplements* (i.e., appositives, parentheticals) are, it seems to me, a validation of the basic tools and techniques of current theoretical linguistics. No major technical move has been made based on supplements, and relatively few linguists have called upon them even as secondary evidence. To a great degree, more obviously integrated constructions bear the burden of supporting current structural and semantic hypotheses.

Yet those hypotheses extend readily to supplements. This is a central claim of Potts 2002a,b, which together form a study of *As*-parentheticals and nonrestrictive relative clauses, henceforth *supplementary relatives*, as well as the relationship between them. The general outlook is this: supplements appear different and untamable, but this is not so when one looks closely. The present chapter pursues this working strategy, refining, expanding, and improving the analyses in Potts 2002a,b. I suggest that a missing piece in those analyses is a precise notion of conventional implicatures (CIs). Though Grice (1975) seems not to have had supplements in mind, their content matches perfectly the definition of CIs that he formulated, repeated in (3.1).

(3.1)
a. CIs are part of the conventional (lexical) meaning of words.
b. CIs are commitments, and thus give rise to entailments.
c. These commitments are made by *the speaker of the utterance* "by virtue of the meaning of" the words he chooses.
d. CIs are logically and compositionally independent of what is "*said* (in the favored sense)", i.e., independent of the at-issue entailments.

The abstract theoretical connection between supplements and CIs is at the heart of this chapter. But the specific analyses are equally important, as supplements, though underexplored, represent a robust and important aspect of natural language. In order to streamline the discussion and avoid overlap with Potts 2002a,b, the present chapter focusses on *nominal appositives* (NAs) such as (3.2a), though *As*-parentheticals and supplementary relatives play a supporting role. Parenthetical adverbials receive a separate treatment, in section 3.7.

(3.2)
a. Ames, a successful spy, is now behind bars. (*nominal appositive*)
b. Ames was, as the press reported, a successful spy. (*As-clause*)
c. Ames, who was a successful spy, is now behind bars. (*supplementary relative*)
d. Amazingly, they refused our offer. (*speaker-oriented adverb*)
e. Thoughtfully, Ed destroyed the evidence for us. (*topic-oriented adverb*)

f. Just between you and me, Aldo is a dangerous spy.
(*utterance-oriented adverb*)

The CI hypothesis for these expressions is extremely well supported. Their content is speaker-oriented and evaluated externally to any intensional or quantificational context, yet a range of evidence indicates that they are syntactically integrated, as adjunct modifiers (section 3.4). The propositions they express are not contextually determined, and the constructions themselves trace back to the conventional meanings of specific lexical items. Moreover, it is easy to imagine the supplement and the main clause having differing semantic values. These properties uniquely identify CI content.

The proposal I defend is that supplements determine routine modifier structures, as in (3.3); the distinction between at-issue and CI content is a semantic one.

(3.3) a. They shot Clyde, who is a wanted fugitive, in the head.

b.

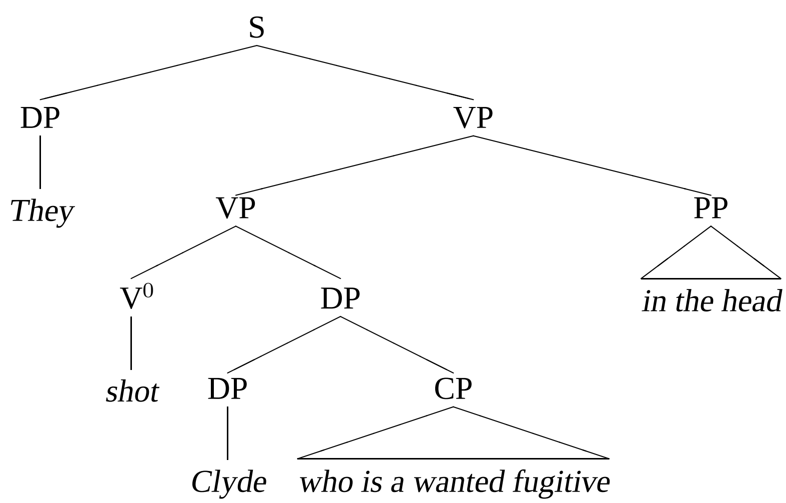

It is the business of chapter 5 to address an alternative, historically prior, analysis of supplements that assigns them a highly nonstandard syntax. I develop this analysis enough to see that it is not a replacement for a theory of CIs. At best, it translates the central features of the CI description logic of chapter 2, $\mathcal{L}_{\text{CI}}$, into tree-structural terms, a project that is doomed to an objectionable level of clumsiness since the notions we seek to capture are semantic.

Once the syntax is fixed as suggested by (3.3), supplements have a great deal to offer the theory of CIs, in large part because of their internal complexity. To date, presumed CI items have been rather limited in their syntactic novelties. But supplements are different. Their internal syntactic structure can be quite complex. In essence, it is possible to take any constituent expressing a declarative at-issue meaning and place it inside a supplement. This reveals new options, provides new angles for the CI supporter to gain a foothold. In some cases, I can sharpen conceptually difficult questions like, 'Is generalization X syntactic, or is it semantic?' At the very least, I am able to show that firm answers to questions like this follow from supportable premises in the analysis of supplements.

The chapter has the following general outline: I first establish some terminology, to facilitate discussion of these relatively unfamiliar constructions. I then present my CI-based analysis in basic form, and move from there to justifying the conservative syntax for these structures that the analysis adopts. The investigation is important in its own right, and also helps to fix the ideal interpreted structure (one with roughly the same structure as the syntactic tree). With the syntax fixed, I discuss the basic semantic properties of NAs, and then show how to apply the CI logic of chapter 2 to this domain. This is as much a case study of NAs as it is an illustration of how $\mathcal{L}_{\text{CI}}$ works. I close by addressing parenthetical adverbs; the discussion of utterance-oriented adverbs, illustrated in (3.2f), provides a bridge between this chapter and chapter 4, since these adverbs have characteristics of both supplements and expressives.

3.2 Some descriptive terminology

The study of supplements in formal linguistics, especially formal semantics, is still young. This afford a rare opportunity to establish a genuinely useful stock of descriptive terms for these constructions. *The Cambridge Grammar of the English Language* (Huddleston and Pullum 2002), the descriptive grammar with the most factually rich discussion of supplements to date, provides terminology that is the product neither of historical accident nor theoretical preconceptions, but rather of clear-sighted descriptive work. I mostly adopt the terms of Huddleston and Pullum's grammar here.

3.2.1 The term 'supplement'

The use of 'supplement' to pick out the class of expressions that includes parentheticals and appositives is due to Huddleston and Pullum 2002. The term is particularly apt in the context of the present work, since CIs as Grice (1975) defined them could be called supplementary semantic content: just as we can slide a newspaper advertising supplement into the trash without loss of journalistic content, so too can we trim a sentence's semantic parsetree of all its terms in τ^c without changing the at-issue content.

I should stress that 'supplement' is merely descriptive terminology. I offer it as a tool for talking informally about *As*-parentheticals, intonationally isolated adverbs, and others. The question of whether 'supplement' should be assigned theoretical content is a more difficult one. I proceed cautiously. The analyses developed in this chapter hinge on a semantically contentful feature COMMA. One could gather together all the constructions that employ a version of this feature under the heading 'supplement'. But it is not clear that this would yield any theoretical dividends. In the worst case, it could result in arbitrarily drawn boundaries, as there are items that meet the intuitive syntactic and intonational conditions for supplementhood, and yet cannot be classified as CI contributors, because they impact the at-issue content. Two prominent examples are the *slifting* construction of Ross (1973) and tag-questions (Culicover 1992), exemplified in (3.4a) and (3.4b), respectively.

(3.4) a. Max, it seems, is a Martian.

b. Max is a Martian, isn't he?

Both of these constructions are intonationally isolated. But neither qualifies as supplementary in the sense that we can remove it and find the at-issue core unaffected by the change. Example (3.4a) is equivalent to *It seems that Max is a Martian*, and (3.4b) questions whether Max is a Martian. There is no sense in which either of them offers the at-issue proposition that Max is a Martian and contributes a secondary proposition qualifying this in some way (cf. *Max, as it seemed, is a Martian*). While there is a good case for analyzing these as syntactically quite like *As*-clauses and clausal supplementary relatives, respectively, the semantic parallels are basically nonexistent. (Though mainly concerned with syntax, Ross (1973:151, fn. 21) can be credited with this insight in the area of slifting and *As*-parentheticals.) I conclude that the term 'supplement' probably does not pick out a semantically significant class of constructions. However, as we will see, the feature COMMA picks out an important subset of supplementary constructions, namely, those with a CI-based semantics.

3.2.2 The pieces of nominal appositives

The bulk of the descriptive work of this chapter is given over to a study of NAs. Since this construction has been looked at only sporadically to date (Barwise and Perry 1983:156–158; Aoun et al. 2001; Elbourne 2001:268–269), it is worth fixing some terms for talking about its three main components, as illustrated in (3.5).

(3.5)

$$\overbrace{\underbrace{\textbf{Chuck}}_{\text{anchor}}\;,\;\underbrace{\textbf{a confirmed psychopath}}_{\text{appositive}}\;,}^{\text{nominal appositive (NA)}}$$

comma intonation

Leftmost is the *anchor* (another term from Huddleston and Pullum 2002). It is (or contains) the syntactic head of the phrase; in section 3.6.4 we will see evidence for this fixed order even in apparently inverted cases like *a cyclist, Lance Armstrong*. Another common term for phrases in this functional position is *associate*, which appears throughout the literature on exceptive phrases such as *no Muppet but Oscar* (Hoeksema 1995; von Fintel 1993; Moltmann 1995). I prefer 'anchor' because it implies no substantive connection with exceptives, and also because it better indicates the role that this part of the NA plays in the overall syntax and in the at-issue semantics.

The intonationally isolated part is the *appositive*. It is set off by the *comma intonation* (Emonds 1976). This intonation is signalled by commas, dashes, or parentheses in print, and by a marked intonation break in speech. On the present analysis, the appositive is generally represented with an at-issue term. In this respect, my proposal revises that of Potts 2002a,b, in which a morpheme internal to the appositive is responsible for its special, non-at-issue semantics. The revisions are motivated in large part by factual considerations pertaining to constructions that contain the same words as these supplements but lack a distinguished comma intonation and (in turn) lack a CI semantics (section 3.6.3). These parallels should be reflected in the grammar. It is therefore reassuring that $\mathcal{L}_{\mathrm{CI}}$ provides a cleaner description if the comma intonation is responsible for the shift into the CI dimension than it does if this shift happens inside the appositive, with the comma intonation left unexplored, as it is in Potts 2002b:650–651.

3.2.3 Relative clause nomenclature

Traditionally, relative clauses are distinguished as 'restrictive' or 'nonrestrictive'. This assumes a model-theoretic perspective on the split. I'd like to promote an alternative: *supplementary* relatives for the semantically nonrestrictive kind, and *integrated* for the potentially restrictive kind. The terminology is from Huddleston and Pullum (2002:§12.4.2, §5.14.3, §15.5.1), where it is observed that 'restrictive' and 'nonrestrictive', read contentfully, divide the class of relatives incorrectly. For instance, not all 'restrictive' relatives are genuinely restrictive:

(3.6) a. the positive numbers that aren't negative
b. the bachelors who are unmarried

The relative clauses in (3.6) do not impose any restrictions that are not already entailed by the head noun. Even if we decided to talk about nonstandard structures where they could be restrictive — models in which we have married bachelors, say — we would not be much better off with regard to the terms 'restrictive' and 'nonrestrictive', because their meanings would in effect contain a hidden modality ranging over possible structures. This level of complexity is highly undesirable in descriptive terminology.

One would like to be able to say that the relatives in (3.6) are nonrestrictive, and yet that they are not 'nonrestrictive relative clauses', i.e., that they are not supplementary relatives: they lack the characteristic comma intonation of supplements; they lack the obligatory *Wh*-pronoun of supplementary relatives; and they are extraposable, (3.7a), whereas supplementary relatives are not, (3.7b).

(3.7) a. The game show host spoke with the bachelors earlier who were unmarried.

b. *The game show host spoke with Eduardo(,) earlier(,) who is unmarried.

Indefinite noun phrases provide another reason not to divide the class of relative clauses along model-theoretic lines. When the anchor (head-noun) of the relative is indefinite, truth-conditions alone do not distinguish the two kinds of relative in monoclausal, extensional environments. The easiest way to see this is to assume that the indefinite contributes a restricted free variable, as in the dynamic semantics of Heim (1982) and Kamp and Reyle (1993) (or the static restatement of Cresswell (2002)). I illustrate in (3.8) with the values we derive in the CI logic (in an intensional model $\mathcal{M}_i$ relative to an assignment g) for typical supplementary and integrated relatives.

(3.8) a. A plumber that endorses phrenology came by.

b. $[\![\lambda w.\, \textbf{plumber}_w(x) \wedge \textbf{endorse}_w(\textbf{phrenology})(x) \wedge \textbf{come-by}_w(x)]\!]^{\mathcal{M}_i,g}$
$= \{w : g(x)$ is a plumber in w and $g(x)$ came by in w and $g(x)$ endorses phrenology in $w\}$

(3.9) a. A plumber, who endorses phrenology, came by.

b. $$\left[\!\!\left[\begin{array}{c} \lambda w.\, \textbf{plumber}_w(x) \wedge \textbf{come-by}_w(x), \\ \lambda w.\, \textbf{endorse}_w(\textbf{phrenology})(x) \end{array} \right]\!\!\right]^{\mathcal{M}_i,g}$$
$$= \left\langle \begin{array}{c} \{w : g(x) \text{ is a plumber in } w \text{ and } g(x) \text{ came by in } w\}, \\ \{w : g(x) \text{ endorses phrenology in } w\} \end{array} \right\rangle$$

We have different semantic objects in the two cases, but the differences are not detectable by intuition. Here is a procedure for seeing that they are identical in all relevant respects: let the set in (3.8) be called W. Form the cross-product $W \times W$. Now we have the same kind of semantic object as we do in (3.9). Finally, assume that a contextual update of a tuple of propositions $\langle \mathbf{p}_1, \ldots, \mathbf{p}_n \rangle$ is a pointwise application of each $\mathbf{p}_i$ to the context. With these assumptions in place, updating with the first member of $W \times W$ is equivalent to the total effect of updating each of the sets in the pair in (3.9). Updating with the second members of $W \times W$ is redundant. We could of course make special assumptions about how we update tuples of propositions, to capture, in a dynamic setting, the secondary nature of supplementary (CI) content relative to at-issue content. But the motivation for this would have to be a great deal more subtle than simply considerations of truth and falsity.

This semantic equivalence indicates that restrictiveness is not the crux of the widely recognized split in the class of relatives. Thus, I adopt the term *supplementary relative* for those with a structure like (3.10) and use *integrated relative* for structures like (3.11).

(3.10)

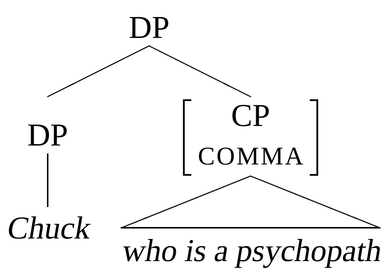

(3.11)

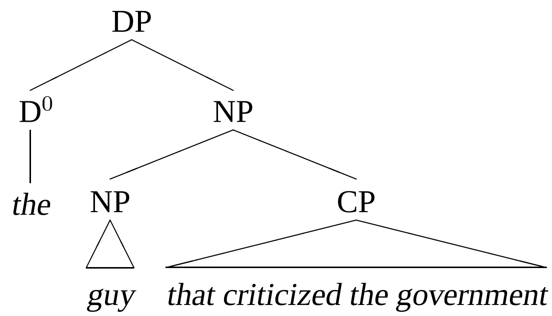

As is evident, the differences are numerous: most prominently, the supplementary relative adjoins to the full nominal in the syntax, whereas the integrated relative adjoins below the determiner. Additionally, the supplement is marked with the feature COMMA, which isolates it intonationally. The semantic translation of this feature delivers CI content. The details are given below, in section 3.6, but I provide a typical example here, so that we can see how 'supplement' might have a semantic interpretation as well. (3.12) is the parsetree associated with (3.10). (For the most part, I limit analyses to extensional types. For these, one can assume that all the predicates' intensional arguments have been quietly filled by a fixed world variable $w_{@}$.)

(3.12)

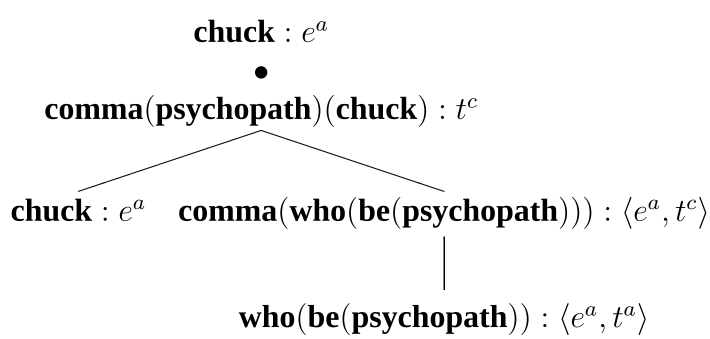

The content of the supplementary relative is independent of the at-issue content. From the point of view of the at-issue semantics, the CI term might as well not be present in the structure. This, again, is a direct formalization of the clause in Grice's (1975) definition that separates the at-issue and CI dimensions, clause (3.1d).

In contrast, the semantic value of (3.11) is given entirely by the lone type e^a at-issue term

$$\iota x[\textbf{guy}(x) \wedge \textbf{criticize}(x, \textbf{the-government})]$$

I extend the term *anchor* to supplementary relatives; the underlined part of the supplementary relative in (3.13) is its anchor.

(3.13) <u>Chuck</u>, who is a confirmed psychopath,

We're now equipped with some language for talking about these constructions. The task is to support an analysis in terms of the CI logic.

3.3 The analysis in brief

A virtue of using the CI logic $\mathcal{L}_{\text{CI}}$ in an analysis of supplements is that it affords a surface-true analysis: syntactically, supplements can appear as regular modifiers. Their CI-based meanings (terms of type τ^c) separate them from the surrounding content, thereby providing a formalization of the intuition that they represent a distinct dimension of meaning.

This chapter explores the factual properties of supplements using $\mathcal{L}_{\text{CI}}$ as a descriptive tool. The task is to understand better the nature of the various parts of supplements, especially NAs, through a description in terms of $\mathcal{L}_{\text{CI}}$. This should in turn lead to an increased understanding of the way $\mathcal{L}_{\text{CI}}$ can be used in natural language analysis. Since I attempt to interweave the descriptive and the formal goals, it is useful to begin by walking through a basic case like (3.14); both the syntactic structures and the semantic parsetrees employed here are motivated in later sections.

(3.14) Lance, a cyclist, is training.

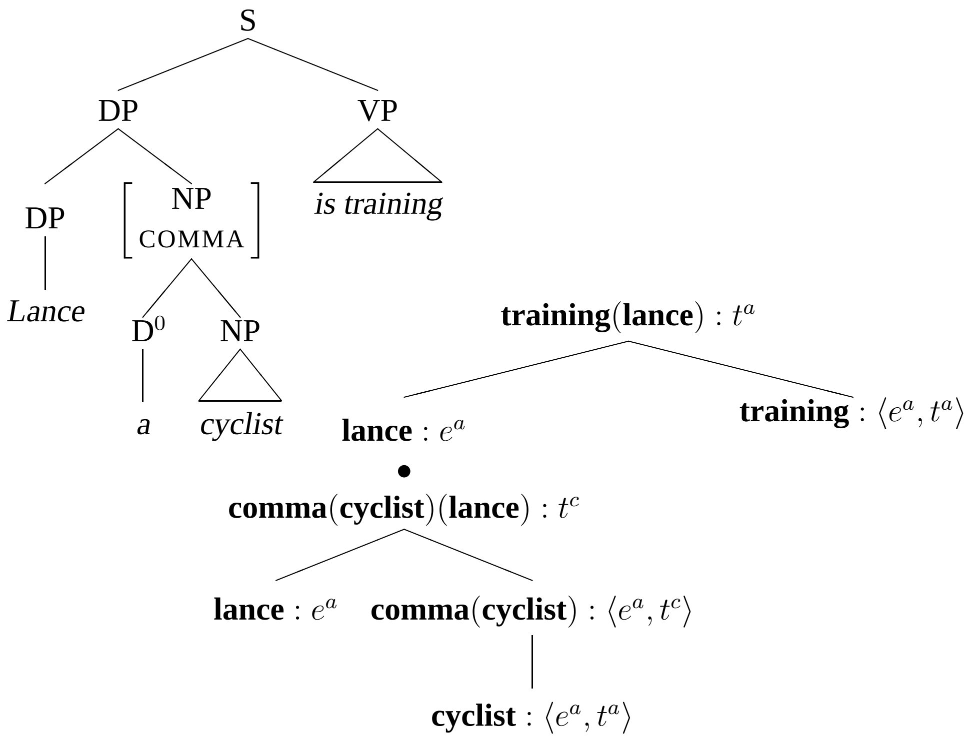

As usual, I represent natural language objects as pairs — a syntactic structure (represented graphically on the left) and a semantic parsetree (on the right). Since the bulk of the formal reconstruction of the at-issue/CI divide rests on the nature of the types, I always provide explicit typing information in the semantic parsetrees. The phrase **lance** : e^a glosses as 'the term **lance** is of the at-issue-entity type', for example.

The analysis centers around the syntactic feature COMMA. It is a signal to isolate the subtree it dominates intonationally, accounting for the commas in print and the intonational boundary marks in speech. Semantically, it performs a type-shift: it takes at-issue content to CI content. For most NAs, the requisite denotation is (3.15), which takes at-issue property-types to CI property-types (watch the change in the superscript on the final t in the input and output types).

(3.15) $\textsc{comma} \rightsquigarrow \lambda f \lambda x.\, f(x) : \langle\langle e^a, t^a\rangle, \langle e^a, t^c\rangle\rangle$

In later sections, I locate reasons to assume that this translation is just one of a handful of the possibilities for COMMA.

I assume that feature terms like **comma** do not introduce terminal elements into either the syntax or the semantics. Rather, I posit a special rule, **feature interpretation**, for including them in the semantic structures:

(3.16) **feature semantics**

$$\begin{array}{c} \beta(\alpha) : \tau \\ | \\ \alpha : \sigma \\ \bullet \\ \boxed{\gamma : \upsilon^c} \end{array}$$

(where β is a designated *feature term* of type $\langle\sigma, \tau\rangle$)

This allows for the introduction of feature terms. (Chapter 2, section 2.6.6, defines a procedure for mapping syntactic nodes with features to nonbranching semantic subparsetrees.) Rule (3.16) licenses the subtree in (3.14) that is of the form

$$\begin{array}{c} \textbf{comma}(\textbf{cyclist}) : \langle e^a, t^c\rangle \\ | \\ \textbf{cyclist} : \langle e^a, t^a\rangle \end{array}$$

The rest of the subject's parsetree is licensed by the central admissability condition of the CI logic, which I repeat in (3.17).

(3.17) **CI application**

$$\beta : \sigma^a$$
$$\bullet$$
$$\alpha(\beta) : \tau^c$$
$$\alpha : \langle \sigma^a, \tau^c \rangle \qquad \beta : \sigma^a$$
$$\bullet \qquad\qquad \bullet$$
$$\gamma : \rho^c \qquad \delta : \upsilon^c$$

The local-tree

$$\textbf{lance} : e^a$$
$$\bullet$$
$$\textbf{cyclist}(\textbf{lance}) : t^c$$
$$\textbf{lance} : e^a \qquad \textbf{comma}(\textbf{cyclist}) : \langle e^a, t^c \rangle$$

has a form specified in (3.17). (Recall that optional material is inside dotted lines and that the linear order of the terminals is irrelevant.)

The composition procedure "leaves behind" the CI material — it is not inherited by dominating nodes or sent directly to a designated point in the parsetree (the root, say). The provision for interpreting such structures that I define in chapter 2 simply folds CI propositions (extensionally, terms in t^c) into the denotation of the entire structure as though they were root-level assertions. I repeat the relevant rule here:

(3.18) **parsetree interpretation**

Let $\mathcal{T}$ be a semantic parsetree with the at-issue term $\alpha : \sigma^a$ on its root node, and distinct terms $\beta_1 : \langle s^a, t^c \rangle, \ldots, \beta_n : \langle s^a, t^c \rangle$ on nodes in it (extensionally, $\beta_1 : t^c, \ldots, \beta_n : t^c$). Then the interpretation of $\mathcal{T}$ is the tuple

$$\left\langle [\![\alpha : \sigma^a]\!]^{\mathcal{M}_i, g}, [\![\beta_1 : \langle s^a, t^c \rangle]\!]^{\mathcal{M}_i, g}, \ldots, [\![\beta_n : \langle s^a, t^c \rangle]\!]^{\mathcal{M}_i, g} \right\rangle$$

where $[\![\cdot]\!]^{\mathcal{M}_i, g}$ is the interpretation function, taking formulae of the meaning language to the interpreted structure $\mathcal{M}_i$, relative to a variable assignment g.

Finally, the highest subtree in (3.14) is licensed by the condition called **at-issue application**, which I repeat here:

(3.19) **at-issue application**

$$\begin{array}{ccc} & \alpha(\beta) : \tau^a & \\ \alpha : \langle \sigma^a, \tau^a \rangle & & \beta : \sigma^a \\ \bullet & & \bullet \\ \boxed{\gamma : \rho^c} & & \boxed{\delta : \upsilon^c} \end{array}$$

One can check that the root local-tree in (3.14) meets this condition. Thus, the entire structure is licensed by the logic.

The basic pieces laid out above can be put together in a multitude of ways, resulting in grandiose objects. For instance, multiple supplements can generally be stacked on a single anchor, as in (3.20).

(3.20)
a. "Colin Powell's son, Michael, Bush's choice to chair the FCC, is an unabashed free-marketeer convinced that Clinton/Gore's procorporate policies on the media were somehow bad for business."[1]

b. "Torriti, shaved, shined, decked out in a tie and sports jacket and a freshly laundered shirt, was uncommonly low-keyed, not to mention sober."[2]

c. The reporter interviewed Lance Armstrong, a rider for the U. S. Postal team, a cancer survivor.

The CI logic assigns these cases a right-adjunction structure. I provide in (3.21) the semantics for (3.20c). (To simplify the diagram, I give the translation of *a cancer survivor* as **survivor** and the translation of *a rider for the U. S. Postal team* as **rider**.)

[1]Miller, Mark Crispin. 2001. *The Bush Dyslexicon*. New York: W. W. Norton and Company (p. 155).
[2]Littell, Robert. 2002. *The Company*. New York: Penguin Books (p. 233).

(3.21)

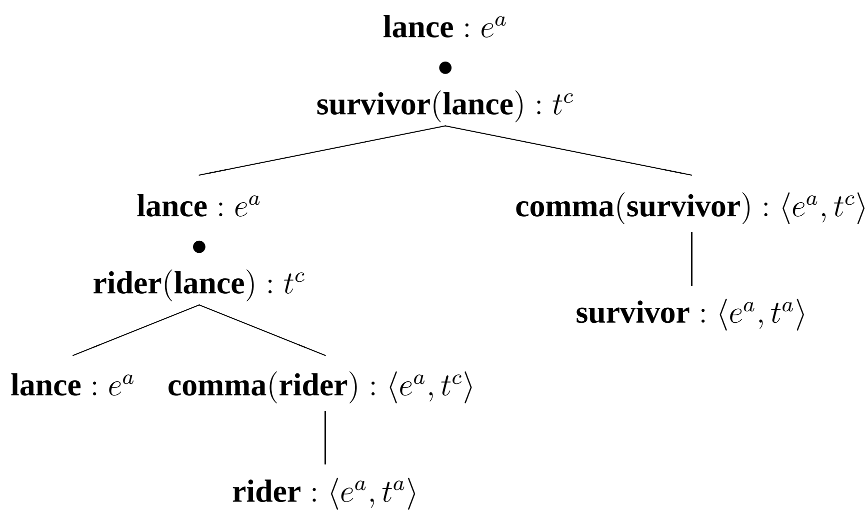

Readers who are familiar with McCawley's (1998:447) analysis of supplementary relatives might think that this is a point of contrast between NAs and supplementary relatives, which McCawley regards as unstackable based on examples like (3.22), which I present with his judgment.

(3.22) ?? Sam Bronowski, who took the qualifying exam, who failed it, wants to retake it.

However, I think that this conclusion is hasty; it seems that supplementary relatives are in general stackable. The problem with (3.22) is probably that it conducts a small narrative using only the supplements. When the content does not build in a story-like fashion, the examples seem unexceptionable. In this, I agree with Kempson (2003a), who offers examples like (3.23a); I include (3.23b, c) to show that supplementary relatives and NAs intermingle quite freely.

(3.23) a. The sole, which I caught yesterday, which was caught in Scotland, was delicious. (Kempson 2003a:(19), with some alterations)

b. I rented *Annie Hall*, which is Woody Allen's finest, a true classic, in order to reminisce about the East Coast U. S.

c. I rented *L.A. Story*, always in demand, which is Steven Martin's finest, in order to reminisce about the West Coast U. S.

But, of course, most of McCawley's (1998) observations are on target. He qualifies his discussion of (3.22) with the following, in which one supplementary relative is buried inside another:

(3.24) Sam Bronowski, who took the qualifying exam, which almost everyone failed, did brilliantly on it. (McCawley 1998:482, fn. 13)

This is also a possibility for NAs. I provide in (3.25) a particularly complex case, which I owe to Helen Majewski (p.c., 2/03).

(3.25) Leo, a lion, a mighty species, swallowed the trainer whole.

Much of the complexity of this examples lies in its dual use of *a lion*. As the argument to *a mighty species*, it is a kind-denoting term: predicates like *species* are defined not for individuals but for natural kinds. But as the functor that applies to the individual Leo, *a lion* is a property.

The analysis of kinds presented by Chierchia (1998) permits this sort of shifting back and forth: we can say that *a lion* denotes the unique plural individual composed of all the lions; the shift from the property to the plural individual denotation is accomplished with the *kind* operator, $^{\cap}$. To return this meaning to the domain of properties, we call upon the *de-kind* operator, $^{\cup}$, which takes entities to properties. These operators are defined in (3.26), in which $\leqslant$ is the part-of relation.

(3.26) a. $^{\cap} = \lambda X.\, \iota x[\forall y[X(y) \leftrightarrow y \leqslant x]] : \langle\langle \sigma^a, t^a\rangle, e^a\rangle$
(*e.g., properties to entities*)

b. $^{\cup} = \lambda x \lambda Y.\, Y \leqslant x : \langle e^a, \langle \sigma^a, t^a\rangle\rangle$
(*e.g., entities to properties*)

If we treat these as feature terms, i.e., the sort of terms that determine unary branching structures, then we obtain the following analysis of (3.25):

(3.27)

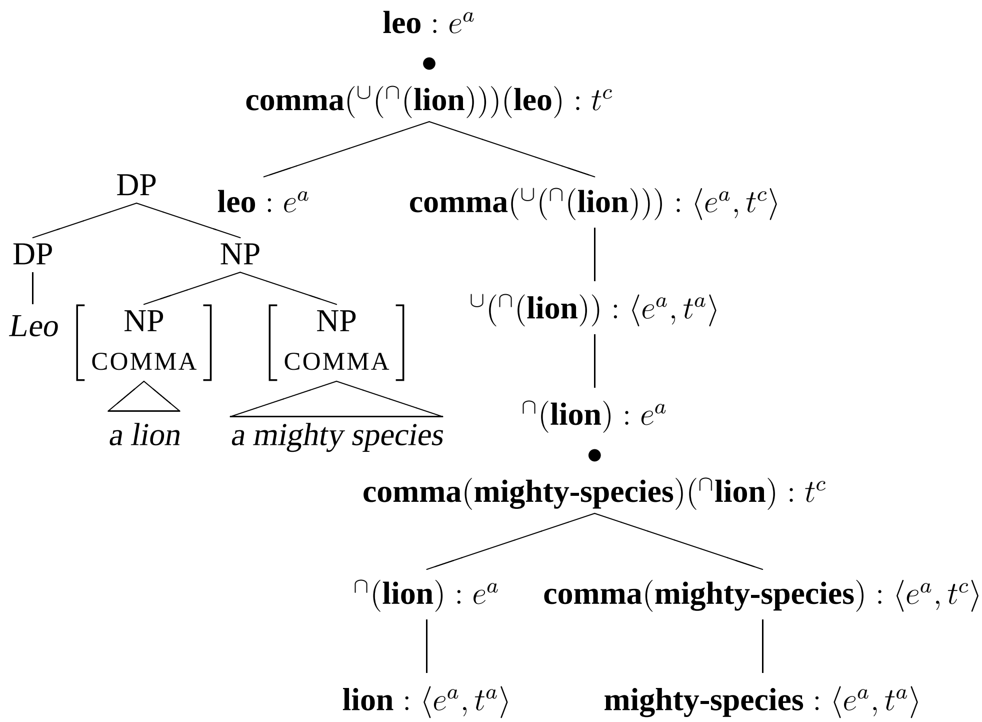

I do not represent *up* or *down* in the syntax, but one could do this for the sake of a transparent mapping. The semantic parsetree is somewhat more complex than the syntax in this case, mainly because the semantics is quite intricate, involving $^{\cap}$**lion** as well as **lion** ($=$ $^{\cup\cap}$**lion**).

Much of the remainder of this chapter is devoted to exploring the ways that the above basic analysis should be generalized to cover a wider-range of NAs, as well as other types of supplements. To ensure that we begin on solid footing, I first motivate the simple adjunction structures for supplements assumed above.

3.4 A conservative syntax

The above description assumes a conservative syntax for supplements, and it begins to make a case for the idea that we can have a clean analysis of the semantics for these expressions without resorting to complex syntactic manipulations to get the right structure for interpretation. Nonetheless, one might think that a nonstandard syntax is desirable for supplements. By and large, the syntactic work that addresses these constructions assigns them a decidedly novel syntax, sometimes shunting them off into a separate structural dimension entirely. Chapter 5 explores this alternative in depth. The purpose of the present section is to argue that narrowly syntactic considerations converge on the conservative, modifier-based syntax suggested by the structures so far. Again, I use mainly NAs to illustrate these claims; Potts 2002a,b do much the same work for supplementary relatives and *As*-parentheticals, respectively.

3.4.1 Adjacency

A generalization that holds across all supplements is that they must be immediately adjacent to whatever constituent they are dependent upon for their interpretation. I exemplify first with NAs:

(3.28) a. *We spoke with Lance before the race, the famous cyclist, about the weather.

b. *Jan was the fastest on the course, the famous German sprinter, yesterday.

c. *Lance has, the famous cyclist, taken the lead.

The clause-final modifiers are present to block a right dislocation or free-adjunct reading of the supplements. Potts 2002b fully documents this restriction for *As*-clauses; here I provide examples involving supplementary relatives.

(3.29) a. *We spoke with Lance before the race, who is a famous cyclist, about the weather.

b. *Jan was the fastest on the course, who is a famous German sprinter, yesterday.

c. *Lance has, who is a famous cyclist, taken the lead.

Similarly, a speaker-oriented adverb can modify only the proposition expressed by the phrase that it is immediately adjacent to; (3.30a) obligatorily involves the meaning of *amazingly* taking the meaning of *it couldn't send a repairman out to help*, whereas (3.30b) can only involve *amazingly* applying to the entire matrix clause.

(3.30) a. The company said that, amazingly, it couldn't send a repairman out to help.

b. Amazingly, the company said that it couldn't send a repairman out to help.

If we analyze these modifiers as adjuncts, then the fact that they are invariably interpreted in their surface positions is a special case of the much broader generalization that (non-lexically-selected) modifiers are always interpreted where they sit at the surface. In a broad class of cases, we can build the requirement into the grammar by make the supplementary expressions semantically unsaturated. I discuss adverbials like (3.30) in greater detail in section 3.7 below, but I can illustrate this point even in the absence of specifics: an adverb like *amazingly* takes propositions into propositions. To capture this descriptive insight, we assign it a lexical entry with the form in (3.31).

(3.31) *amazingly* $\rightsquigarrow \lambda \mathbf{p} \lambda w.\, \mathbf{amazing}_w(\mathbf{p}) : \langle\langle s^a, t^a\rangle, \langle s^a, t^c\rangle\rangle$

This entry narrows down the class of parsetrees that the meaning of *amazingly* can appear in to just those where it is sister to a term α of type $\langle s^a, t^a\rangle$. The value of their mother will be $\lambda w.\, \mathbf{amazingly}_w(\alpha)$.

One might think that extraposition is a potential source of incorrect predictions, but in fact the CI logic neatly explains why CI-contributing supplements are not extraposable. Consider the structure of an NA extraposed from subject position:

(3.32)

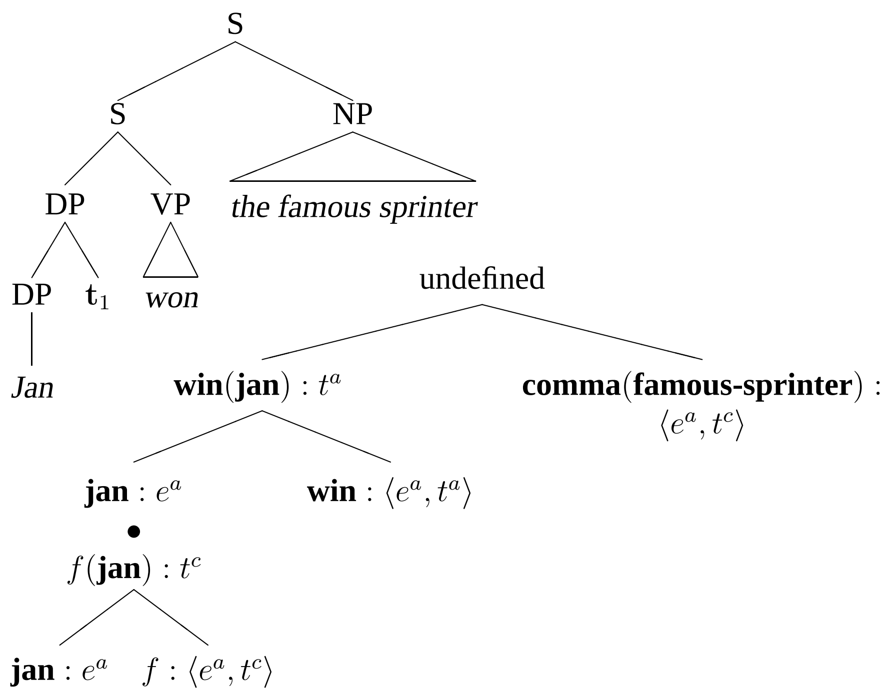

There is no provision for abstracting over the free CI functional variable f so that its value can be given by **famous-sprinter**. Moreover, even if we were to formulate a principle that would permit abstraction over the highest occurrence of the variable f, the requisite expression would have to be of type $\langle\langle e^a, t^c\rangle, t^c\rangle$. But this is a type whose first member is a CI type. We have no such types in the space to work with provided by the type definition, which I repeat here:

(3.33)
a. e^a, t^a, and s^a are basic at-issue types for $\mathcal{L}_{\text{CI}}$.
b. e^c, t^c, and s^c are basic CI types for $\mathcal{L}_{\text{CI}}$.
c. If σ and τ are at-issue types for $\mathcal{L}_{\text{CI}}$, then $\langle\sigma, \tau\rangle$ is an at-issue type for $\mathcal{L}_{\text{CI}}$.
d. If σ is an at-issue type for $\mathcal{L}_{\text{CI}}$ and τ is a CI type for $\mathcal{L}_{\text{CI}}$, then $\langle\sigma, \tau\rangle$ is a CI type for $\mathcal{L}_{\text{CI}}$.
e. If σ and τ are at-issue types for $\mathcal{L}_{\text{CI}}$, then $\langle\sigma \times \tau\rangle$ is a product type for $\mathcal{L}_{\text{CI}}$, a subset of the set of at-issue types for $\mathcal{L}_{\text{CI}}$.
f. The full set of types for $\mathcal{L}_{\text{CI}}$ is the union of the at-issue and CI types for $\mathcal{L}_{\text{CI}}$.

Thus, nothing extra need be said to block extraposition. Various factors conspire to make it impossible within the bounds of the CI logic. The deepest reason is that the sister of the extraposed item in the semantic parsetree must be decorated with a term that takes CI types into CI types, a possibility that is disallowed by the logic.

It is worth contrasting this account of the nonextraposablity of supplements with the condition Emonds (1979) offers, which is simply that a supplement and its anchor must be immediately adjacent to each other. This kind of ad hoc statement does not formalize the restriction (indeed, formalization would be extremely challenging given the structures that Emonds works with), whereas it is built into $\mathcal{L}_{\text{CI}}$. What's more, Emonds' analysis wrongly disallows all stacking of supplements (McCawley 1998:453); we saw in (3.23) that this is overly stringent: stacking is possible, a fact that the $\mathcal{L}_{\text{CI}}$ treatment captures without extra statement.

3.4.2 Right-adjunction

Another argument for the right-adjunction view of NAs in particular derives from the apparent absence of such clauses in languages that forbid right-adjunction categorically. Turkish, for instance, seems not to have NAs; the closest construction is an integrated left-adjoining version that is close to the English construction *The bicyclist Hasan.* An example is given in (3.34). (My thanks to Jorge Hankamer for data and discussion.)

(3.34) Un-lU bisiklet-Ci Hasan-la yarIS-tan Once konuS-tu-k.
*fame-ous bicycle-ist Hasan-with race.*ABL *before speak.*PAST*-we*
'We spoke with Hasan, the famous bicyclist, before the race.'

Turkish seems also not to have syntactically, morphologically, or intonationally distinguished supplementary relative clauses. I believe that these considerations extend to Japanese as well. The idea is that right-adjunction provides the basis for an explanatory account, in the form of a simple deduction: NAs are right-adjoining; Turkish does not permit right-adjunction; therefore, Turkish does not permit NAs. Though this is not in itself an explanation, it highlights an important correlation.

3.4.3 Case-marking

Case-marking in German provides a compelling argument for treating NAs as nominals containing a modifier structure, as in (3.14) above. In NAs, the anchor and the appositive share case, which is whatever case is normally determined by the predicate that takes the construction as its argument (Durrell 1995:§2.6). For example, in (3.35a), the NA appears as the object of the preposition *mit* ('with'), which governs the dative case for its object. Though the anchor *Jan*, a proper noun, does not show morphological case, such case shows up on the appositive, in both the article and the adjective. I include (3.35b) to indicate that case is in fact marked on both the anchor and the appositive where neither is a proper name.

(3.35) a. Wir sprachen mit Jan, dem weltberühmten Radfahrer.
*we spoke with Jan the.*DAT *world-famous.*DAT *cyclist.*
'We spoke with Jan, the world famous cyclist.'

b. Ich sah meinen Freund, den Pfarrer.
*I saw my.*ACC *friend, the.*ACC *parson*
'I saw my friend, the parson.' (Durrell 1995:37)

c. nach dem Todes meines Onkels, des früheren Bürgermeisters der Stadt
*after the death my.*GEN *uncle.*GEN*, the.*GEN *former.*GEN *mayor.*GEN *the.*GEN *city*
'after the death of my uncle, the former major of the city,'
(Durrell 1995:37)

Durrell (1995:37–38) notes some exceptions. A genitive anchor usually takes a nominative or dative appositive; a weekday given with the preposition *am* ('on.DAT') can take a dative or accusative appositive date (e.g., *am Montag,* {*den*/*dem*} *16. Juni* ('on Monday, the 16th of July')); and the genitive is common after a prepositional phrase headed by *von.* These details seem idiosyncratic; I won't attempt to account for them. But one class of exceptions, not noted by Durrell, suggests that some, but not all, NAs are in fact reduced supplementary relatives or topicalized structures. In general, the case-marking facts provide an argument that NAs do not reduce to supplementary relatives. If that were the correct analysis, one would expect to find nominative case on the appositive, since this is the case of both arguments to the copula in German:

(3.36) a. Der Kerl ist ein weltberühmter Radfahrer.
*the.*NOM *guy is a.*NOM *world-famous.*NOM *cyclist*
'The guy is a world famous cyclist.'

b. Wir sprachen mit Jan, der ein weltberühmter
*we spoke with Jan, who.*NOM *a.*NOM *world-famous.*NOM
Radfahrer ist.
cyclist is
'We spoke with Jan, who is a world famous cyclist.'

An analysis that assigned the NA *Jan, dem weltberühmten Radfahrer* a structure in which *welberühmten Radfahrer* were a post-copular argument would get the case-marking facts wrong unless a special stipulation were made that ellipsis altered the case-marking conventions.

But for some apparent NAs, nominative case is what one finds:

(3.37) "Sie steigen wieder in den Wagen, ein grosser schwarzer, [. . .]"
*they climb again into the.*ACC *car.*ACC*, a.*NOM *big.*NOM *black.*NOM*,* [. . .]
'They climbed back into the car, a big black one, . . .'[3]

We get the case marking facts right if we assume that this is in fact a reduced supplementary relative clause. It seems inappropriate to claim that the structures for these cases involve ellipsis of the form ~~der~~ *ein grosser schwarzer* ~~ist~~, as this is a nonconstituent deletion. It seems possible that the structure instead involves fronting of the predicative nominal to a topic phrase above CP, with subsequent deletion of that CP complement to the topic head, as proposed for fragment answers by Merchant (2003). The result of such an analysis of (3.37) is as in (3.38) (strikeout on a node u indicates that the subtree rooted at u is phonetically empty; 'Top' abbreviates 'Topic'):

[3] Leonard, Elmore. 2000. *Schnappt Chili*. Translated by Hans M. Herzog. München: Wilhelm Goldmann Verlag (p. 126).

(3.38)

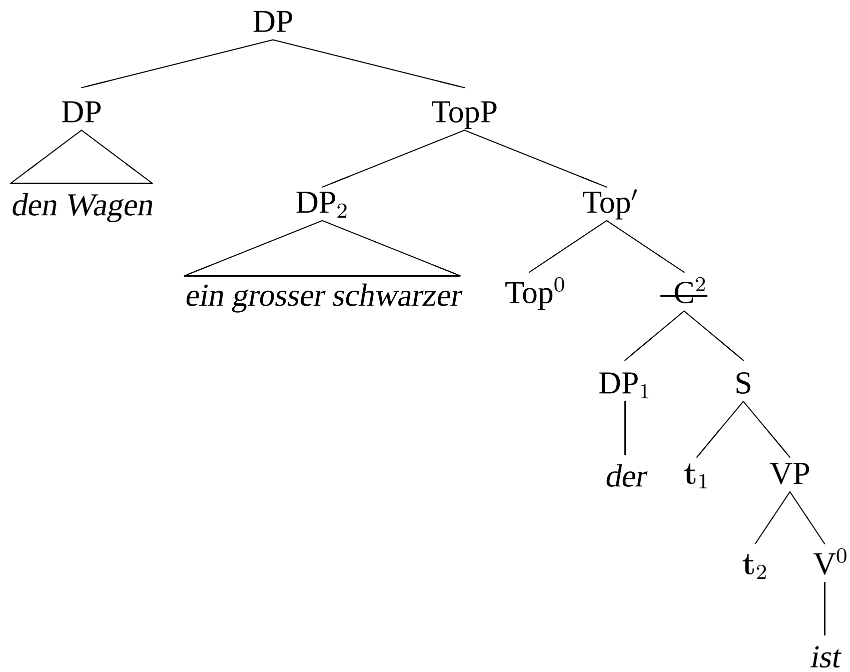

For genuine NAs, the structure involves mere adjunction of an NP:

(3.39)

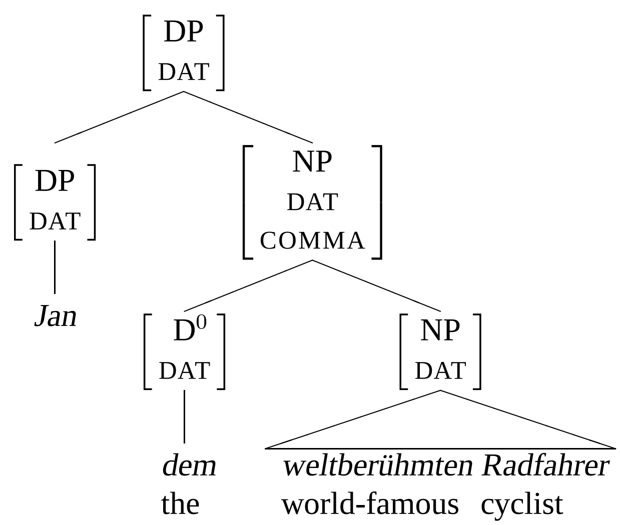

Here we can take advantage of the presence of COMMA to state the case-marking facts. We simply need to say something equivalent to the following: if an NP with the feature COMMA is adjoined to a DP, then the case-marking features of DP appear on NP.

The case for treating nominative-marked appositives as reduced supplementary relatives is supported by certain contrasts involving definiteness marking. As reviewed below, NAs differ from supplementary relatives in the area of definiteness marking:

(3.40) a. Lance Armstrong, the cyclist, is training now with his cycling buddies.

b. #Lance Armstrong, who is the cyclist, is training now with his cycling buddies.

c. My idol, the cyclist, is in Texas. #My idol, the astronaut, is in Ohio.

The expected uniqueness effects associated with the definite article are not in force in (3.40a). They return, though, in the supplementary relative counterpart, (3.40b). I include (3.40c) to show that the NA does not function as a complex definite description in these cases; the anchor has all the usual properties when it is definite.

The same contrasts hold for German:

(3.41) a. Jan, der Radprofi, hat 1997 die Tour gewonnen.
*Jan the.*NOM *cycling-pro has 1997 the Tour won*
'Jan, the cycling pro, won the Tour in 1997.'

b. #Jan, der der Radprofi ist, hat 1997 die Tour gewonnen.
*Jan who.*NOM *the.*NOM *cycling-pro is has 1997 the Tour won*
'Jan, who is the cycling pro, won the Tour in 1997.'

Thus, the claim that nominative-marked appositives are reduced supplementary relatives when their anchors have non-nominative case has some bite in this area: if the claim is correct, then we should see supplementary relative-like uniqueness presuppositions akin to (3.41b). This seems to be correct:

(3.42) a. Sie sprachen mit Jan — immerhin ein mehrfacher Tour
*they spoke with Jan after-all a.*NOM *many-time.*NOM *Tour*
Sieger!
winner
'They spoke with Jan, after all a many-time Tour winner.'

b. #Sie sprachen mit Jan — immerhin der mehrfacher Tour
*they spoke with Jan after-all the.*NOM *many-time.*NOM *Tour*
Sieger!
winner
'They spoke with Jan, after all the many time Tour winner.'

We must make an exceptions for epithets such as *das Arschloch* in example (2.30), repeated here.

(3.43) Sie haben ja den Webster – das Arschloch — endlich gefeuert.
they have JA the Webster the asshole finally fired
'They JA finally fired Webster, the asshole.' (Kratzer 1999)

Here, we find no uniqueness effects; the presence of multiple individuals deserving the predicate named by *Arschloch* does not lead to infelicity. But this seems to be a general fact about epithets: their definite-marking is not to be taken seriously. The most direct English translation of the above NA is probably *Webster, that asshole*, which similarly lacks a uniqueness presupposition. When we leave the realm of expressive content, though, the correlation between nonmatching nominative case and definiteness presuppositions seems strong.

I offer these examples merely to show that, in both German and English, NAs contrast with supplementary relatives in terms of their internal syntax and semantics. Externally, though, they seem to require the same general treatment.

3.5 Basic semantic properties

All the tests outlined in chapter 1 for classifying semantic content as conventionally implicated converge on the hypothesis that supplements have a CI-based semantics.

3.5.1 Nondeniable meanings

It is easy to dispel the premise that their content is conversationally implicated. The best argument for this position is that it is not contextually dependent, and hence it is not deniable, (3.44a), nor is it suspendible with epistemic riders, (3.44b, c).

(3.44) a. Edna, a fearless leader, started the descent. #Edna is not a fearless leader.

b. #Lance Armstrong, the 2002 Tour winner, is training, if Armstrong did win the 2002 Tour.

c. #If Armstrong did win the 2002 Tour, then Lance Armstrong, the 2002 Tour winner, is training.

These facts also suggest that a presuppositional treatment is not feasible. Example (3.44c) is especially useful in this regard: if the proposition that Armstrong is the 2002 Tour winner were a presupposition engendered by the NA in the consequent, then the preposed *if*-clause would work to satisfy its requirements, that is, the example would work in the same fashion as the classic example in (3.45).

(3.45) If Eddie has a dog, then his dog is a ferocious man-eater.

The example does not assert or presuppose that Eddie has a dog. But (3.44c) does assert that Armstrong is the 2002 Tour winner, hence the oddness of placing this content inside the antecedent of a conditional.

3.5.2 Antibackgrounding

Recall that Grice (1975) makes no provision that CI meaning should be entailed by the discourse participants' common ground at the point of utterance (van der Sandt 1988:74); definition (3.1) does not address this question, whereas it is of course at the heart of all presupposition logics that presuppositional content is backgrounded in this way. We can use this fact to build a rather strong case against treating supplements as presupposed. In general, supplements are used to introduce new information, often as a means of further clarifying and contextualizing the at-issue content around them. As such, if their content is backgrounded, then the result is infelicity due to redundancy. Let's call this the *antibackgrounding effect*. It is illustrated in (3.46), in which the initial sentence sets up a context for the NA-containing sentence (3.46a) as well as (3.46b), which contains the factive predicate *know*.

(3.46) Lance Armstrong survived cancer.

a. #When reporters interview Lance, a cancer survivor, he often talks about the disease.

b. And most riders know that Lance Armstrong is a cancer survivor.

Example (3.46b) is completely felicitous. The initial sentence satisfies the presuppositional requirements of *know* by placing into the common ground the content of the complement to *know*. If we think of presuppositions in the basic terms laid out by van der Sandt (1992), the presupposition in (3.46b) finds an appropriate antecedent in the initial sentence.

In contrast, when the NA's content finds an antecedent — when its content is already entailed by the common ground — the NA is infelicitous due to redundancy. It would take rather special discourse conditions to license the NA in (3.46a).

It is worth lingering over the logic of this argument. Some recent uses of the term "presupposition" seem to have drastically relaxed the requirement that presupposed content be entailed by the context prior to the point of utterance. For many authors, presupposed content that does not meet this condition can be accommodated without conscious effort, which basically places it outside the bounds of detection by linguistic techniques.

For instance, Steedman (2000:654) writes that, for a certain range of cases, "the listener rapidly and unconsciously adjusts his or her model of the domain of discourse to support the presuppositions of the speaker". This considerably reduces the distance between presuppositions and at-issue entailments in a way that might prevent us from obtaining a unified theory of one or both of them. But even if we grant the revision, it remains the case that presuppositions *can* be backgrounded at any time. This alone never leads to redundancy. But the antibackgrounding requirement for supplements does not merely say that NAs and other expressions need not be backgrounded, it says that they *cannot* be backgrounded.

3.5.3 Nonrestrictiveness

Though I retreated from the terms 'restrictive' and 'nonrestrictive' for the two basic types of relative clause, the notion of restrictiveness is significant for a theory of supplements. Though it is false to say that restrictive relatives are always genuinely restrictive, it remains true that supplements cannot restrict the at-issue value of their anchors. Perhaps the strongest indication that this is the case derives from contrasts between

NAs and integrated appositive expressions like (3.47b) (Jackendoff 1984; Huddleston and Pullum 2002:§16).

(3.47) a. Armstrong, the Texan, is a cyclist. #Armstrong, the Ohioan, is an astronaut.
(NAs)

b. Armstrong the Texan is a cyclist. Armstrong the Ohioan is an astronaut.
(integrated appositives)

The integrated supplement in (3.47b) is restrictive. We seem to have here an instance in which a name like *Armstrong* can be used descriptively, rather than as a term of direct reference. So we have evidence that this reading of the name exists. This possibility still does not license (3.47a). It seems that here we must fix a referent for *Armstrong*. The appositive then functions to assert that its denotation has the property of being a Texan. The second sentence is infelicitous given the first and the assumption that Armstrong is not both a Texan and an Ohioan. The bottom line is that we cannot use the NA to restrict the value of its anchor. The integrated appositive does allow this restriction, however.

Similar contrasts are easily constructed with pairs of supplementary and integrated relatives. The CI logic captures this nonrestrictive property of supplements, and indeed of all CI modifiers, by completely separating them from the at-issue content. To make things concrete, I offer an analysis of the integrated appositive in (3.49), which is centered around the functional, descriptive meaning for *Armstrong* in (3.48).

(3.48) *Armstrong* $\rightsquigarrow \lambda f.\, \iota x[\textbf{called}(\text{Armstrong})(x) \wedge f(x)]$

(3.49)

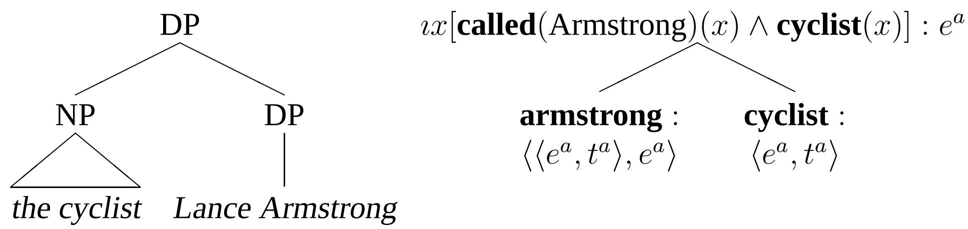

The defining feature of the analysis is that the proper name *Armstrong* is realized as a function from predicates to definite descriptions (entities).

3.5.4 Independence of meaning

In chapter 1, section 1.4.3.1, I used NAs to exemplify the independence of CIs from the at-issue content, thereby validating clause (3.1d) of Grice's (1975) definition. I repeat the relevant example in (3.50a), along with additional supplements that show the same thing.

(3.50) a. Lance Armstrong, an Arkansan, has won the 2002 Tour de France!
b. Lance Armstrong, who is an Arkansan, won the 2002 Tour de France.
c. Bill Clinton is from Arkansas, as is Lance Armstrong.

As you probably know, Lance Armstrong is a Texan, not an Arkansan. But he is the 2002 Tour de France winner, and Bill Clinton is from Arkansas. The fact that the values differ, and, more to the point, the fact that the CI is false, does not prevent us from recovering information from these sentences in the present context. This is contrary to the guiding idea that presuppositions are conditions on definedness for the at-issue content (i.e., where they are false or undefined, the at-issue content itself is undefined). For CIs, however, we seek dimensional independence.

3.5.5 Widest-scope interpretation

As noted above, there is a persistent intuition that supplements are not syntactically integrated, but rather adjoined to the root, possibly via a relation that is disjoint from dominance (Lakoff 1966; McCawley 1982, 1987, 1989, 1998; Huddleston and Pullum 2002; Emonds 1976; Culicover 1992). The primary motivation for this analysis does not derive from narrowly syntactic facts. As discussed in section 3.4 above, there probably are not narrowly syntactic arguments for a nonstandard syntax. Issues of constituency, case-marking, and language-specific structural restrictions point to a modifier-structure for those supplements under discussion here. The feature of supplements that these root-level adjunction analyses seek to capture is semantic: supplements are always interpreted as though they took widest (primary) scope.

To some, this suggests a presuppositional treatment (Keenan 1971). However, Beaver (2001) raises this possibility, in connection with his example (3.51), and then proceeds to dismantle it.

(3.51) "Sweden may export synthetic wolf urine — sprayed along roads to keep elk away — to Kuwait for use against camels."[4]

Beaver concludes on the basis of such examples that the presupposition holes (negation, modalization, questioning, conditionalization) do not provide the means for obtaining a sufficient condition for presuppositionhood. They indicate that certain kinds of content can take scope outside of a hole and yet not count as presupposed. He writes that the proposition that wolf urine is sprayed along the roads to keep elk away

[4]Associated Press, January 19, 1995 (cited in Beaver 2001:20, (E34)).

(3.52) “appears within the scope of the modal operator ‘**may**’ [...But] one would infer [...] that synthetic wolf urine is sprayed along roads to keep elks away. Thus, according to the embedding-under-modals test (and others can be applied with the same result) this inference should be designated as presuppositional. But many theories associate presuppositions with information which is in the common ground between interlocutors, or assumed by the speaker to be in this common ground. On such an account the anti-elk application of synthetic wolf urine would not appropriately be termed presuppositional, since the writer of the text very likely does not expect readers to have any previous knowledge of the subject.”

(Beaver 2001:20)

Beaver’s example also militates against a treatment in terms of at-issue entailments, since these certainly do embed under presupposition holes like *may*.

Example (3.53) further supports this claim that NAs do not embed.

(3.53) It’s false that Alonzo, a big-shot executive, is now behind bars.

The matrix *it’s false that* negates the proposition that Alonzo is now behind bars. But the NA’s meaning (that Alonzo is a big-shot executive) survives; there is no reading of this example on which it means

$$\neg\left(\begin{array}{r}\textbf{behind-bars}(\textbf{alonzo}) \quad \wedge \\ \textbf{big-shot-executive}(\textbf{alonzo}) \quad\;\;\end{array}\right)$$

This analysis would make (3.53) true if Alonzo were a lowly intern but not jailed.

In building a case against a presuppositional treatment of a given meaning, it is wise to check the properties of that meaning when the expression that gives rise to it is embedded below a presupposition plug like *say* or *believe*, since presuppositions are generally interpreted in the scope of these operators (though definite descriptions often scope outside of them). This test further supports a nonpresuppositional account of NAs; see (3.54).

(3.54) a. The agency interviewed Chuck, a confirmed psychopath, just after his release from prison.

b. $\approx$ Chuck is a confirmed psychopath, and the agency interviewed Chuck just after his release from prison.

There does seem to be a genuine semantic parallel here, one that a theory of supplements should capture. One might infer from these cases that the interpreted structure for (3.54a) is in fact (3.54b).

However, this analysis is of limited utility. We can see this quite clearly when we move to considering syntactically embedded NAs like that in (3.55).

(3.55) a. Sheila believes that the agency interviewed Chuck, a confirmed psychopath, just after his release from prison.

b. $\not\approx$ Sheila believes that Chuck is a confirmed psychopath and that the agency interviewed Chuck just after his release from prison.

Here, we do not attribute to Sheila the proposition that Chuck is a confirmed psychopath, though the material expressing this is in the clausal complement of *believe*. Once again, the NA receives widest scope. What this means in terms of building an analysis based on coordination is that we cannot simply take the coordinate-like structure that we arrived at for (3.54a) and assume that it can be embedded compositionally under, say *believe*. We must in fact deem such structures ungrammatical when embedded. There are technical tricks for ensuring this widest-scope property, tricks we can accomplish in the semantics or the syntax (see chapter 5). But it seems that we miss something central about these constructions in advancing this stipulation and not pressing it any further.

I offer (3.56) as an attested instance in which allowing embedded readings would create what is quite clearly an unwanted ambiguity.

(3.56) "Consider the following situation. In front of Ralph stand two women. For some reason we don't need to investigate, Ralph believes that the woman on the left, who is smiling, is Bea, and the woman on the right, who is frowning, is Ann. As a matter of fact, exactly the opposite is the case. Bea is frowning on the right and Ann is smiling on the left."[5]

The sentence of interest is the second. The entire point of the puzzle is to fix the position of the smilers and the frowners *in the actual world*. But the clauses expressing these facts appear embedded inside the complement to *believe.* If there are embedded readings of supplementary appositives, then the sentence in question is ambiguous between a reading that attributes the supplement content only to Ralph and a reading in which it is interpreted in the actual world. It seems safe to say that the embedded reading is blocked here. This is why it works to say that what Ralph believes is "exactly the opposite" of what actually obtains.

The CI logic is capable of ensuring the widest-scope effects for these modifiers. It does this via the usual mechanisms: the logic quite handily ensures that CI content — and by entailment, NA content — never ends up in the scope of anything.

3.5.5.1 Apparent exceptions

Some speaker feel that the claim that supplements are semantically unembeddable is too strong. These intuitions do not jibe with those advanced in the literature, nor do

[5] Aloni, Maria. 2000. *Quantification under Conceptual Covers*. Ph.D. thesis, Universiteit van Amsterdam. Published in the ILLC Dissertation Series, 2001–1 (p. 65).

they comport with the way that supplements are actually used. But it is important to inquire into the source of such intuitions.

To my knowledge, the only published claims for embedded readings of supplements appear in Boër and Lycan (1976), during their discussion of supplementary relatives. They claim to find a "dialect difference" concerning examples like (3.57).

(3.57) It's false that Dick, who is an expert on Austin, loves the Bonzo Dog Band.
(Boër and Lycan 1976:20, (35))

(3.58) It is not the case that Dick, who is an expert on Austin, loves the Bonzo Dog Band.
(Boër and Lycan 1976:20, (36))

Speakers of what Boër and Lycan (1976) call Dialect A treat these examples as semantically equivalent to (3.59), which is roughly the result we derive in the CI logic but with a classical conjunction where we derive an ordered pair.

(3.59) $\textbf{expert-on}(\textbf{austin})(\textbf{dick}) \wedge \neg\textbf{love}(\textbf{bonzo-dog-band})(\textbf{dick})$

Speaker of Dialect B "freely grant that (3.59) is far and away the more natural reading of (3.58)" (p. 20), but they also say that (3.59) can be read with the negation taking scope over the supplementary relative and the main clause, producing something equivalent to (3.60).

(3.60) $\neg\begin{pmatrix} \textbf{expert-on}(\textbf{austin})(\textbf{dick}) & \wedge \\ \textbf{love}(\textbf{bonzo-dog-band})(\textbf{dick}) & \end{pmatrix}$

On this reading, (3.58) is *true* if Dick loves the Bonzo Dog Band but is not an expert on Austin.

Boër and Lycan (1976) go on to say that negation is the only operator that permits this kind of ambiguity for any speakers. They provide examples involving modal operators (p. 22, (47)) and propositional attitude predicates like *be convinced* (p. 21, (45)), and say for each that the embedded reading of the supplementary relative is unavailable. They then discuss the conditions under which Dialect B speakers can embed a supplementary relative. They turn out to be rather special circumstances: special "stage-setting" is required. And they say, quite tellingly, that the only way to be sure that the supplement ends up in the scope of negation is to place it inside a quotation; (3.61) is their illustrative example:

(3.61) It is simply false that "Dick, who is an expert on Austin, loves the Bonzo Dog Band." (Boër and Lycan 1976:23, (49))

It seems clear that Boër and Lycan (1976) do not provide counterexamples to the claim that supplements, at least of the relative kind, are always evaluated as root-level assertions, i.e., that they never end up in the scope of anything. We can instead argue that the special stage-setting and highly specific intonation contour required to derive (3.60)-like readings amount to direction quotation. The question then arises: 'How would we go about trying to determine whether genuinely embedded readings exist?'

One must keep in mind that nothing about placing a supplement inside a propositional attitude context will entail that the subject of that propositional attitude verb *disbelieves* that supplement's content. In (3.62), we do not attribute to Sheila the proposition that Chuck is *not* a psychopath, nor do we even indicate that Sheila is agnostic about the truth of this proposition.

(3.62) Sheila believes that Chuck, a psychopath, is fit to watch the kids.

What's more, propositional attitude verbs like *believe* do not impose any exhaustivity requirements, even as conversational implicatures. In saying "Sheila believes that

Chuck is fit to watch the kids" one does not suggest that this is the only thing Sheila believes. So one will never arrive at the implicature that the subject of a propositional attitude verb does not endorse the content of an embedded supplement using only simple cases like (3.62).

This seems an obvious point. But it points up the limited claim I am making about supplements — they that are evaluated as main clause assertions no matter where they appear. One can syntactically embed them, but from the point of view of the semantic environment in which they appear, it is as though they were not even there.

In light of this situation, I see just one way to test for cases of semantic embedding: follow the sentence in question with an explicit disavowal, in a main clause utterance, of the content of the supplement. Such a continuation should reduce possible readings of the supplement to an embedded reading if such exists, since the primary-scope reading is inconsistent. A relevant test case:

(3.63) Sheila believes that Chuck, a confirmed psychopath, should be locked up. #But Chuck isn't a confirmed psychopath.

I marked this example according to my intuitions, which are reflected in the analyses cited above by McCawley, Emonds, and others. They strike me as compelling evidence that we do not have genuinely embedded readings. I venture that apparently embedded readings, including those of Dialect B, involve direct quotation. (Chapter 4 contains a similar argument involving expressives.)

3.5.5.2 Discourse subjunctive

An interaction between the German discourse subjunctive, Konjunktiv I, and clausal supplements seems at first to constitute an exception to the usual speaker-orientation of supplement content. I discuss the Konjunktiv I at length in chapter 4, section 4.7, so here I provide only a brief description. In general, Konjunktiv I marking on the inflected auxiliary of a clause C indicates that the speaker is not publicly committed to the proposition expressed by C. A simple example:

(3.64) Juan behauptet, dass Maria krank sei.
*Juan maintains that Maria sick be.*KONJ
'Juan maintains that Maria is sick.'

The subjunctive marking is a kind of explicit disavowal of commitment. Thus one might expect it to be impossible inside supplements if they are speaker-oriented. In fact, Konjunktiv I combines fairly readily with an embedded clause:

(3.65) Juan behauptet, dass Maria, die sehr schwach sei, krank sei.
*Juan maintains that Maria who very weak be.*KONJ *sick be.*KONJ
'Juan maintains that Maria, who is supposed to be really weak, is sick.'

It seems at first that we have contradictory demands: the relative clause commits the speaker to the proposition that Maria is sick, whereas the Konjunktiv I registers a denial of commitment. However, when one studies the distribution of Konjunktiv I more broadly, one finds that it can occur in main clauses provided that the context includes an agent to whom the content of the clause can be relativized. For example:

(3.66) Juan behauptet, dass Maria krank sei. Sie sei sehr schwach.
*Juan maintains that Maria sick be.*KONJ *She be.*KONJ *very weak*
'Juan maintains that Maria is sick. According to him, she is very weak.'

The effect is a kind of modal subordination (Schlenker 2003; von Stechow 2002). I do not, in this work, attempt to understand how this subordination occurs, but it is very clear that syntactic main clauses can be semantically embedded when they bear the Konjunktiv I marking. (This can happen to a limited degree in English as well, but speakers run the risk of ambiguity, since there is no systematic way to indicate noncommitment.) Since supplements and main clauses are interpreted in essentially the same fashion in $\mathcal{L}_{\mathrm{CI}}$, the fact that Konjunktiv I is possible in both is not a surprise after all. The prediction $\mathcal{L}_{\mathrm{CI}}$ makes is that, for example, (3.65) does commit the speaker to the proposition that Juan (or perhaps some other discourse-salient agent) maintains that Maria is very weak. The prediction is correct.

3.5.6 Definites and indefinites

3.5.6.1 Definites in NAs and supplementary relatives

The examples in (3.67) indicate an intriguing contrast between NAs on the one hand and supplementary relatives and main clauses on the other.

(3.67) a. Armstrong, the cyclist, is training with his cycling buddies.
b. # Armstrong, who is the cyclist, is training with his cycling buddies.
c. # Armstrong is the cyclist.

As discussed briefly above in connection with German supplements, it seems as though the usual uniqueness presuppositions associated with the definite article are called off when it appears as the syntactic head of a predicative NA. Example (3.67a) is felicitous in a situation with multiple cyclists, whereas (3.67b) is not. This observation appears also in Elbourne 2001:268–269. Elbourne observes that (3.68a) is not at all like (3.68b).

(3.68) a. Amnesty International, the human rights group, claims in a new report that . . .

b. Amnesty International, the only entity which is a human rights group, claims that . . .

Example (3.68b) is likely to be infelicitous in most situations, whereas (3.68a) is fine, especially if the speaker feels that the hearer needs to be reminded of the most salient property of Amnesty International, viz., that it is a human rights group.

Elbourne (2001) chooses not to venture a solution; he writes "I will not pursue this line of inquiry further in the present paper. (The semantics of the definite article on which it relies is the focus of work in progress.)" (p. 269). It seems to me that question turns on the discourse conditions in which NAs appear. I offer the following generalizations:

(3.69) a. If a speaker chooses a definite article to head an NA's appositive, then the proposition expressed by that NA is deemed essential by the speaker to determining the referent of the anchor.

b. If a speaker chooses an indefinite article to head an NA's appositive, then the proposition expressed by that NA is deemed essential by the speaker to the narrative.

The basic idea is that a speaker's use of *Lance Armstrong, the cyclist* indicates that the property of being a cyclist helps the hearer to better understand what entity *Lance Armstrong* picks out. It is possible that the information that Lance is a cyclist plays no other role in the current narrative. For instance, in a story about famous people who have survived cancer, *Lance Armstrong, the cyclist* might appear, even though Armstrong's cycling has little if anything to do with his bout with cancer. In this case, the appositive helps the reader to understand something about the individual under discussion. It is unlikely that *Lance Armstrong, a cyclist* would be used in this context, since cycling is not directly relevant to the narrative.

In contrast, if we find the indefinite article, then the NA's content is essential to the narrative. In the above context, *Lance Armstrong, a survivor of testicular cancer* is far more likely than the version with the definite article, because the appositive provides essential information about why Armstrong is relevant to the story. Indeed, in this situation, my sense is that *Lance Armstrong, the survivor of testicular cancer* would be distinctly odd, perhaps even presupposing that he is the only such survivor (a return of the usual definiteness effects for *the*).

The conditions in (3.69) predict a kind of free variation when the NA expresses a proposition that is essential to the narrative and says something essential about the anchor's denotation. It is a challenge to test these generalizations extensively, since one must provide a considerable amount of background information for each example and then ask about how that example contributes to the context. I offer (3.70) along with an assertion that the definite-headed appositive expresses a defining feature of Torriti as well as an essential piece of the story it is part of.

(3.70) "Torriti, the tradecraft shaman capable of blending into a nonexistent crowd, shed the lazy pose of a fat man who drowned his sense of doom and gloom in PX booze and swung into action."[6]

I recommend that fans of espionage fiction do the necessary work of reading the preceding 206 pages in order to verify my claim. They will not be disappointed. I predict that they will also find that the sentence would have been equally good with an indefinite article.

[6]Littell, Robert. 2002. *The Company*. New York: Penguin Books (p. 207).

We can begin to make sense of why NAs, but not supplementary relatives, are subject to the pragmatic generalizations in (3.69) by inspecting their respective parsetrees. In a supplementary relative, the functor that applies to the post-copular predicate is at-issue; the shift to CI content happens at the top of the relative clause. In an NA, the corresponding predicate is the argument to the functor that takes us to CI content. Evidently, the semantics of the definite article is different in the CI domain than it is in the at-issue domain.

3.5.6.2 A link with specificational clauses

I turn now to a second kind of contrast concerning articles, this one observed by McCawley (1998:468) in relation to (3.71).

(3.71) a. A recent winner of the Illinois State Lottery, Albert Swenson, has announced that he plans to move to Bermuda.
(McCawley 1998:468, (2b))

b. ?? A recent winner of the Illinois State Lottery, who is Albert Swenson, has announced that he plans to move to Bermuda.
(McCawley 1998:468, (2b′))

I believe that the restriction in (3.71b), though certainly real, is not absolute. The supplementary relative *who is Albert Swenson* is essentially a specificational copular clause of the sort studied by Mikkelsen (2002b,c,a, 2003) under the heading *non-clefted specificational clause* (NCSC). Mikkelsen observes that such clauses are generally marked when they have indefinite subjects, but that this is not an absolute restriction. Thus:

(3.72) a. $^{\#}$A recent winner of the Illinois State Lottery is Albert Swenson.

b. $^{\#}$A doctor is John. (Heycock and Kroch 1999:(52b))

c. [You are telling me about your sick pet. I reply]
A vet who might be able to help you is John Roberts.
(Mikkelsen 2002a)

Thus, as far as the study of supplementary relatives is concerned, we just need to ensure that examples like (3.71b) have a composition scheme that mirrors that of NCSCs. The following is a representation that follows closely Mikkelsen's (2002c) analysis:

(3.73)

$\textbf{comma}(\textbf{who}(\textbf{albert})) : \langle e^a, t^c \rangle$

$\textbf{who}(\textbf{albert}) : \langle e^a, t^a \rangle$

$\textbf{who} : \langle e^a, \langle e^a, t^a \rangle\rangle$ $\quad$ $\textbf{albert} : e^a$

$\lambda x.\, x : \langle e^a, e^a \rangle$ $\quad$ $\textbf{albert} : e^a$

This parallels Mikkelsen's (2002c) analysis in that the phrase *is Albert Swenson* has an entity-level denotation. The subject takes this entity as an argument. There is just a single difference between regular NCSCs and the supplementary relative counterpart: according to Mikkelsen, the pre-copular phrase in an NCSC is of type $\langle e^a, t^a \rangle$. In my analysis, the relative pronoun has a slightly different type: it takes entity inputs and returns property-denoting outputs. This is necessary to ensure that the supplementary relative combines properly with the at-issue subject, which must be an individual in this case. But the analysis retains Mikkelsen's (2002c) function–argument structures, as well as most of the denotational assignments she makes. Hence, we can expect her analysis of the semantics and pragmatics of NCSCs to carry over easily to the supplementary domain, thereby reducing the puzzle McCawley poses with (3.71) to that of (3.72).

3.6 The internal structure of NAs

In section 3.3, I provided the guiding ideas for how to apply the CI logic to this domain. It's time now to look at the internal structure of supplements, especially NAs, to see how to capture the conditions on these constructions.

3.6.1 The anchor

In general, quantified expressions are not possible anchors in NAs:

(3.74) a. *Every climber, {an/the} experienced adventurer, was found sipping hot cocoa in the lodge.

b. *No climber, {an/the} experienced adventurer, was found sipping hot cocoa in the lodge.

This is part of an important broader generalization:

(3.75) Nonrestrictive modifiers associate only with referring expressions.
(Thorne 1972:553; Karttunen 1976:367; McCawley 1998:451;
Potts 2002a:83; Huddleston and Pullum 2002:1060; and others)

(3.76) a. *The doctor gave a lollipop to each child, who she examined.
(McCawley 1998:451, (24a′))

b. *Susan interviewed every senator, who is crooked.
(McCawley 1998:451, (24b′))

c. *No person, who knows everything, is perfect.

(McCawley 1998:451, (24c′))

d. *No candidate, who scored 40% or more, ever failed.

(Huddleston and Pullum 2002:1060, (7i))

The CI logic (considered so as to include lexical items) cannot derive meanings for structures in which the anchor is quantified. I illustrate in (3.77).

(3.77)

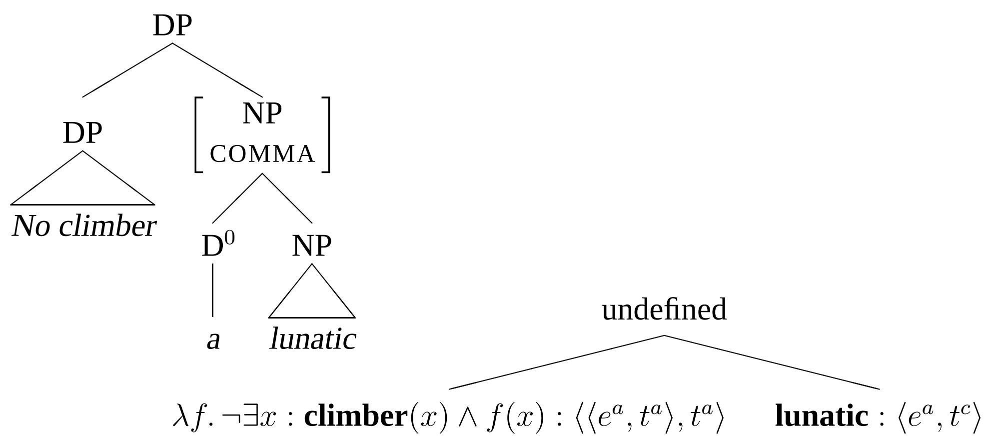

The quantifier cannot take the appositive as its argument, because the appositive is of type $\langle e^a, t^c\rangle$ but the quantifier takes only meanings of type $\langle e^a, t^a\rangle$. So we have a type-mismatch. To ensure that this type-mismatch holds for all structures, we need to say that the appositive cannot shift to type $\langle\langle\langle e^a, t^a\rangle, t^a\rangle, t^c\rangle$, taking the quantifier meaning as its argument. This is a mathematical possibility, but evidently not a linguistic one. The meaning we would derive for *No climber, a lunatic, survived* would be equivalent to *No climber is a lunatic and no climber survived.*

But, for the usual reasons, we needn't worry about such type-shifting. Type-shifting functions are of course terms of our logic. In order to take a CI term τ of type $\langle e^a, t^c\rangle$

into a term of type $\langle\langle\langle e^a, t^a\rangle, t^a\rangle, t^c\rangle$, we would require a type-shifter with the type

$$\langle\langle e^a, t^c\rangle, \langle\langle\langle e^a, t^a\rangle, t^a\rangle, t^c\rangle\rangle$$

But both members of this type are, of necessity, CI types. The type definition does not contain such types. Therefore, we block this otherwise worrisome derivation at the most basic level. The only remaining concern is that the type shift might happen prior to the move to CI types. To block such a composition scheme, we just need to limit the range of types that **comma** can have. I assume that the only possibility is (3.114) below.

So the logic itself handles a complete ban on quantified anchors quite nicely. But the underlying generalization is more complicated than this. An example of a grammatical anchor that has the form of a quantifier is (3.78a), which I owe to Lyn Frazier (p.c., 2/03).

(3.78) a. Every climber, all experienced adventurers, made it to the summit.
b. Every climber, experienced adventurers all, made it to the summit.
c. *Every climber, experienced adventurers, made it to the summit.

How are these cases best handled? In locating an answer, I think it is worth looking at examples like (3.79).

(3.79) The students, most of them linguists, missed the bus.

In this example, the appositive is semantically complete. It is a small-clause of some kind. It needn't apply to the meaning of *the students*, though this appears to be its syntactic anchor. I suggest approaching these examples as follows:

(3.80)

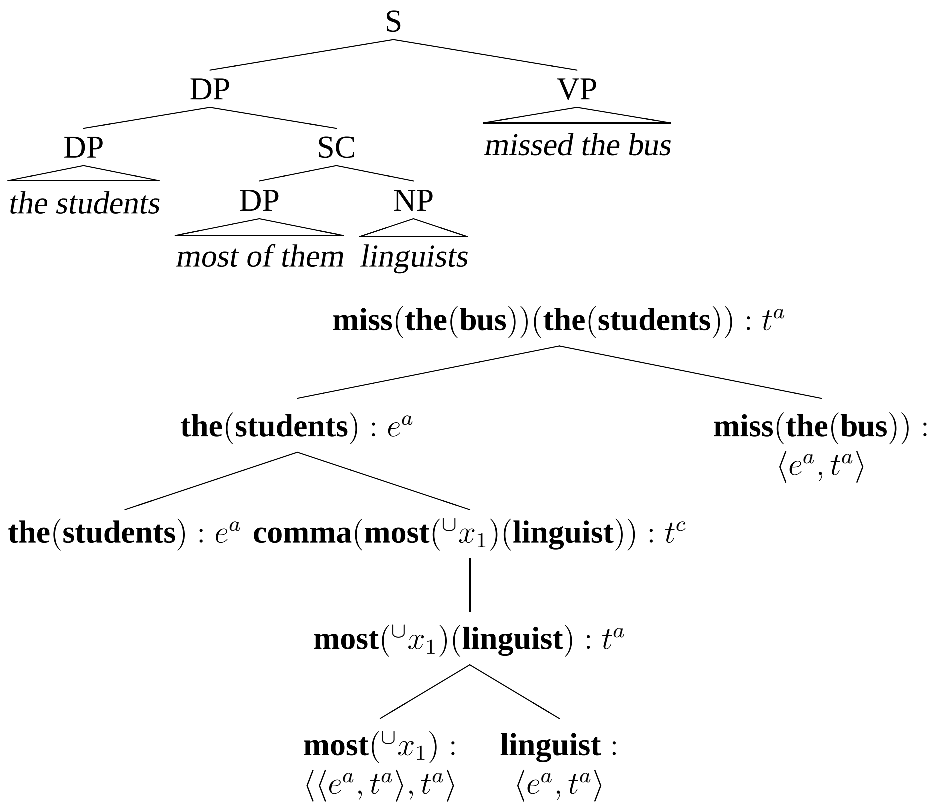

The variable x_1 should have an entity-level denotation that is identical to the denotation of **the**(**students**). The term $^\cup x_1$ provides the required predicate-level denotation for combination with **most**. (The de-kind functor $^\cup$ is defined in (3.26b) and in appendix A.)

To license this structure, we need to call upon the following rule, which is used in a more restricted form in Potts 2003a and has a more general instantiation in chapter 4.

(3.81) **isolated CIs**

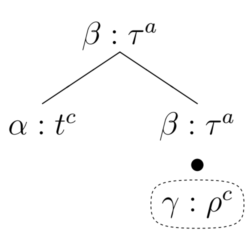

This rule licenses structures in which the CI bears no semantic function–argument relation to its anchor. It is useful for small-clause constructions like (3.79), as well as niched conjunctions such as (3.82a) and interjections like (3.82b).

(3.82) a. Luke has — and you'll never believe this — eaten fifty eggs.

b. Jeb, that guy Dick's always talking about, will someday rule the earth.

And it is also, I claim, useful for analyzing the problematic quantified cases cited in (3.78). An important feature of those cases is that they have an obligatory quantifier of some sort in them. To reduce these cases to examples like (3.79) and (3.82) we need only assume that the restriction on the appositive-internal quantifier is inherited from the restriction on the at-issue determiner. That is, we require structures like the following:

(3.83) a. every climber, adventurers all,

b.

$$\textbf{every}(\textbf{climber}_1) : \langle\langle e^a, t^a\rangle, t^a\rangle$$

$$\textbf{every}(\textbf{climber}_1) : \langle\langle e^a, t^a\rangle, t^a\rangle \qquad \textbf{comma}(\forall y : f_1(y) \rightarrow \textbf{adventurer}(y)) : t^c$$

$$\forall y : f_1(y) \rightarrow \textbf{adventurer}(y) : t^a$$

$$\lambda g.\, \forall y : f_1(y) \rightarrow g(y) : \langle\langle e^a, t^a\rangle, t^a\rangle \qquad \textbf{adventurer} : \langle e^a, t^a\rangle$$

For simplicity, I assume that the restriction on **all** is supplied by a free-variable; the idea is that $[\![\textbf{climber}]\!]^{\mathcal{M}_i,g} = g(f_1) =$ the characteristic function of the set of climbers. (The restriction logic of Gawron (1996) can derive this result by simply reusing the same variable in the quantifications signalled by *every* and *all*.) On this analysis, the universally quantified cases with quantifiers inside their appositive parts are not exceptions to the generalization that the anchor cannot be quantified.

However, it is not clear that this logic extends to all cases in which something that appears quantified occupies the anchor position. Recently, Kempson (2003b) has claimed that even *most* is viable in this position, at least with NRRs. She offers examples like the following:

(3.84) $^{\%}$Most older people on the march, who left after Jesse Jackson, got home without too much trouble.

There is certainly a restrictive reading of this example, represented in (3.85a). There is also, it seems, a reading on which the relative clause picks up only on the restriction on *most*, delivering a reading like (3.85b). What is in dispute is the question of whether there is a reading in which the relative clause ends up in the nuclear scope of the quantifier *most*, as in (3.85c)

(3.85) $^{\%}$Most elderly, who heard Jackson, got home early.

a. restrictive relative reading

$$\{x \mid \textbf{elderly}(x) \wedge \textbf{hear}(\textbf{jackson})(x) \wedge \textbf{get-home-early}(x)\} > \{x \mid \textbf{elderly}(x) \wedge \textbf{hear}(\textbf{jackson})(x) \wedge \neg(\textbf{get-home-early}(x))\}$$

b. universal reading of the relative

$$\left\langle \begin{array}{c} \{x \mid \textbf{elderly}(x) \wedge \textbf{get-home-early}(x)\} > \\ \{x \mid \textbf{elderly}(x) \wedge \neg(\textbf{get-home-early}(x))\}, \\ \forall x[\textbf{elderly}(x) \rightarrow \textbf{hear}(\textbf{jackson})(x)] \end{array} \right\rangle$$

c. relative material in the nuclear scope

$$\{x : \textbf{elderly}(x) \wedge \textbf{hear}(\textbf{jackson})(x) \wedge \textbf{get-home-early}(x)\} > \{x : \textbf{elderly}(x) \wedge \neg(\textbf{hear}(\textbf{jackson})(x) \wedge \textbf{get-home-early}(x))\}$$

Reading (3.85a) simply involves a restrictive relative, and so can be set aside. Reading (3.85b) seems to be of a sort that occurs often with epithets: we have a dependency between the domain of the relative clause and the domain on the at-issue quantifier. Again, the restriction logic of Gawron (1996) is ideally suited to obtaining these readings based on structures in which we appear to have a free-variable in the translation of the relative.

So, from the present perspective, only (3.85c) poses a special difficulty. It places the relative clause in the scope of the quantifier. I am highly skeptical that such readings actually exist. If they do arise, then I conjecture that they do not truly involve supplementary relatives. I see two ways to support this independently. The first builds on observations dating at least to Thorne (1972), who notes that a host of utterance-modifying adverbs can appear inside supplementary relatives, despite the fact that such adverbs resist true syntactic embedding. Some examples to illustrate:

(3.86) a. Incidentally, Jed is a lunatic.

b. *They believe that, incidentally, Jed is a lunatic

c. Jed, who, incidentally, is a lunatic, will be in power one day.

It is also true that supplementary relatives do not force performatives to lose their performative status, though regular syntactic embedding does. One can see this best using the adverb *hereby*, which actually delivers ungrammaticality, not just a loss of performative status:

(3.87) a. I hereby swear off unmotivated bombing campaigns.

b. *George believes that I hereby swear off unmotivated bombing campaigns.

c. I, who hereby swear off unmotivated bombing campaigns, am going to run for office.

A generalization is in the offing: if C is a supplementary relative, then C permits these utterance-level modifiers. With this premise fixed, we can try to mix them with quantifier examples like the above:

(3.88) a. *Most older people on the march, who, incidentally, left after Jesse Jackson, got home without too much trouble.

b. *Every parrot sang a song, which, $\{$between you and me/confidentially$\}$, it didn't understand.

We can also test intuitions about truth and falsity, though these are bound to be more delicate. As observed in chapter 1, and in section 3.5.4 of the present chapter, speakers easily tease apart the values of appositive material from the content of the at-issue dimension; if a speaker reports *Lance Armstrong, an Arkansan, won the 2002 Tour de France*, we can use the value $\langle 1, 0 \rangle$ to characterize the truth value of this sentence in our world (since Armstrong is a Texan, not an Arkansan, but he did win the 2002 Tour). It seems to me that this independence does not extend to cases in which there is purportedly a binding relationship between the two dimensions. I am not an ideal speaker to test this, since I am skeptical that we actually find supplementary relatives and other supplements in the semantic scope of quantifiers, so I just leave this as a potential area in which to research, relying on the clearer adverb-interpolation examples above to justify the decision to treat apparent cases of quantifier–appositive links as manifesting different constructions than the ones analyzed here.

I close this section by addressing a slightly different kind of support for generalization (3.75). Although individual concepts are suitable anchors, (3.89), no anchor can contain a pronoun that is bound from outside of that anchor, as seen in (3.90).

(3.89) The U.S. president, a major force in world politics, changes every eight years at least (we hope).

(3.90) a. Every student_1 spoke with a psychiatrist of hers_1 that welcomes housecalls.

b. *Every student_1 spoke with a psychiatrist of hers_1, a caring individual who welcomes housecalls.

c. Sally_1 spoke with a psychiatrist of hers_1, a caring individual who welcomes housecalls.

Modification of a bound pronoun by an integrated relative is fine, as in (3.90a). And if the antecedent of a pronoun is a referential expression, the result is again fine, as in (3.90b). But an anchor that is or contains a bound variable is impossible, as seen in (3.90c). On the present view, this is entirely expected: in (3.90c), the CI expressed by the appositive is the proposition that $g(x_1)$ is a caring individual who welcomes housecalls. But nothing about the sentence can make $g(x_1)$ a salient individual, since x_1 does not remain free in the at-issue assertion:

(3.91)

$$\textbf{psychiatrist-of}(x_1) : e^a$$
$$\bullet$$
$$\textbf{caring}(\textbf{psychiatrist-of}(x_1)) \wedge$$
$$\textbf{welcome}(\textbf{housecalls})(\textbf{psychiatrist-of}(x_1)) : t^c$$

$$\textbf{psychiatrist-of}(x_1) : e^a \qquad \textbf{comma}\left(\begin{array}{c}\lambda x.\,\textbf{caring}(x) \wedge \\ \textbf{welcome}(\textbf{housecalls})(x)\end{array}\right) : \langle e^a, t^c \rangle$$

$$\lambda x.\,\textbf{caring}(x) \wedge \textbf{welcome}(\textbf{housecalls}\,)(x) : \langle e^a, t^a \rangle$$

The variable in the CI term on the root node remains free. While this structure can appear in a meaningful structure in which the x_1 in the at-issue dimension is bound by the universal, this requires the pronoun *her* to have both a referential and a bound variable semantics. Though nothing about the $\mathcal{L}_{\text{CI}}$ treatment deems this incoherent, it violates fairly fundamental conditions on how ambiguous or homophonous expressions behave. The facts represented by (3.92) have been known since at least Lakoff 1970.

(3.92) a. Ed lies, and Sue does too.

b. Ed and Sue lie.

$=$ Ed speaks untruthfully, and Sue speaks untruthfully.
$=$ Ed gets prone, and Sue gets prone.
$\neq$ Ed speaks untruthfully, and Sue gets prone.
$\neq$ Ed gets prone, and Sue speaks untruthfully.

The mixed readings are not possible. To ensure that we reduce the problem posed by (3.90b) to the one posed by (3.92), we could partition the set of variables into a set $V_b = \{x_1^b, x_2^b, \ldots\}$ of dependent variables and a set $V_f = \{x_1^f, x_2^f, \ldots\}$ of free variables, roughly as in Farkas 1997. This would force a choice in (3.91), and the only available reading would be one on which *her* translates as a free pronoun. For concreteness, assume that we have the following:

(3.93) a. $\textit{lie} \rightsquigarrow \begin{cases} \textbf{speak-untruthfully} : \langle e^a, t^a \rangle \\ \textbf{get-prone} : \langle e^a, t^a \rangle \end{cases}$

b. $\textit{her} \rightsquigarrow \begin{cases} x_1^b : e^a \\ x_1^f : e^a \end{cases}$

So the distinctions are independently motivated. The important point is that the variable in the CI dimension bears no relation to the value of the variable in the at-issue dimension. Hence, the multidimensionality of the theory is essential to the explanation.

3.6.2 The appositive

In most cases, the appositive is property-denoting; the result is that NAs are strongly reminiscent of predicative copular clauses with individual-denoting subjects. The syntax seems not to impose further limitations: nominal, adjectival, and prepositional phrases are all possible:

(3.94) a. The agency interviewed Chuck, a confirmed psychopath, just after his release from prison.

b. Yewberry jelly, toxic in the extreme, will give you an awful stomachache.

c. Ed, in trouble with the law once again, has altered his identity.

In general, quantified appositives are ungrammatical (excepting cases like (3.79) and (3.78), discussed above as involving saturated supplementary meanings). Some examples:

(3.95) a. *We spoke with Tanya, Ashley, and Connie, every secretary in the department, about the broken printer.

b. *We approached Tanya and Ashley, most secretaries in our department, about the broken printer.

c. Armin, Jaye, and Junko, (*all) the phonologists at UCSC, attended the conference.

However, quantifiers that can appear in predicative positions (Partee 1987) are also fine in NAs:

(3.96) a. We spoke with Hillary, no amateur climber, about the dangers.

b. Ed’s house, at one time every color of the rainbow, now has aluminum siding.

c. We spoke with Tanya, everything to everyone around here, about the broken printer.

The parallel between predicative copular constructions and NAs is grounded in the logic itself. In general, theories of predicative copular constructions converge on a function–argument structure of the form represented by (3.97), the parsetree for *Julio is Swedish.*

(3.97)

$$\begin{array}{ccc} & \mathbf{swedish}(\mathbf{julio}) : t^a & \\ \mathbf{julio} : e^a & & \mathbf{be}(\mathbf{swedish}) : \langle e^a, t^a \rangle \\ & \mathbf{be} : \langle \langle e^a, t^a \rangle, \langle e^a, t^a \rangle \rangle & \mathbf{swedish} : \langle e^a, t^a \rangle \end{array}$$

So we have an instance in which the predicate applies to the subject. In NAs, the appositive applies to the anchor. So I offer the following descriptive claim:

(3.98) An expression E can appear as the predicate in a predicative copular construction if and only if E can appear in an NA’s appositive position.

The claim has real bite given the many ways we have of isolating predicative constructions from other kinds of copular clause; I refer to Higgins 1973, Sharvit 1999, and Mikkelsen 2002b,c for various tests and references.

Example (3.99) seems to falsify the right-to-left direction of (3.98); I thank Youri Zabbal (p.c., 2/03) for (3.99a).

(3.99) a. Ed saw three paintings, each more beautiful than the last.
b. *Three paintings were each more beautiful than the last.

Barry Schein (p.c., 2/03) observes that the acceptability of (3.99a) derives from the fact that the supplement is actually predicated of the seeing events described by the main clause. That is, this is not an NA in the sense that *three paintings* is the anchor. We want an analysis that treats these on par with *Ed saw three paintings, saw that each was more beautiful than the last he looked at.* McCawley (1998:468, (5)) observes that this kind of apposition is attested elsewhere; in his (3.100a), the anchor for *a factor...* is the full main clause; I include (3.100b) to show that verb phrases are also fair targets.

(3.100) a. Mercantile's growth is far more broadly based than before, a factor which has enabled the group to live with high interest rates and still keep a firm grip on margins.
b. Tami bicycles in the Santa Cruz mountains, a favorite activity of Brian's as well.

Thus, we see that NAs are but one proper subtype of bare appositive expressions with propositional denotations.

After addressing the nature of the feature COMMA that appears in NAs and other supplements, I return to the question of whether *A former linguist, Edward Witten* and similar examples involve an inverted NA.

3.6.3 Comma intonation

Throughout this chapter, the shift from at-issue to CI content has been achieved by the semantic reflex of the syntactic feature COMMA. The guiding idea is that it is COMMA that enables supplements to meet the specification in Grice's (1975) definition that CIs be lexical.

As observed in Potts 2002b:650, it is sometimes the case that intonation is the only thing that distinguishes a supplement from a fully integrated phrase. This is in effect observed also by Huddleston and Pullum (2002:1065), who cite examples like (3.101).

(3.101) a. She had two sons who were studying law at the university.

b. She had two sons, who were studying law at the university.

In the end, the two examples are true in exactly the same situations, since the relative clauses associate with the same indefinite in both cases (see section 3.2 for discussion). But we find striking differences in their intonational properties. When we embed them, of course, further differences appear:

(3.102) a. It's just false that she had two sons who were studying law at the university. Neither of her sons had been past junior high.

b. It's just false that she had two sons, who were studying law at the university. #Neither of her sons had been past junior high.

In short, the restrictive relative is inside the scope of the negation, whereas the supplementary relative is not.

My CI analysis traces all the semantic differences between these two kinds of relative clause back to the comma intonation. Both relative clause constructions involve the translation in (3.103).

(3.103) *who were studying law at the university* $\rightsquigarrow$
$\lambda x.\,$ **at-the-university**$($**study**$(x)) : \langle e^a, t^a \rangle$

For the integrated case, this meaning combines via intersection with the common noun meaning **son**, which in turn serves as the argument to the cardinal determiner. For the supplementary case, this meaning shifts via **comma** to become a property-denoting CI term:

(3.104) *, who were studying law at the university,* $\rightsquigarrow$
comma$(\lambda x.\,$ **at-the-university**$($**study**$(x))) : \langle e^a, t^c \rangle$

This shift forces adjunction to the phrase *two sons* in the manner described above. Everything flows from the semantics of the single reliable perceptible difference between the two kinds of relative clause.

We can assume that the same differences separate NAs from integrated appositives:

(3.105) a. The visiting scholar, Alonzo, is snoozing in my office. #The visiting scholar, Johan, is snoozing in the conference room.

b. The visiting scholar Alonzo is snoozing in my office. The cyclist Alonzo is cruising around in the mountains.

We'd like the meaning of **visiting scholar** to combine with the meaning of **Alonzo** for (3.105b). For this, I assume (as with example (3.49) above) that *Alonzo* translates as $\lambda f.\, \imath x[$**called**$(\text{Alonzo})(x) \wedge f(x)]$.

(3.106)

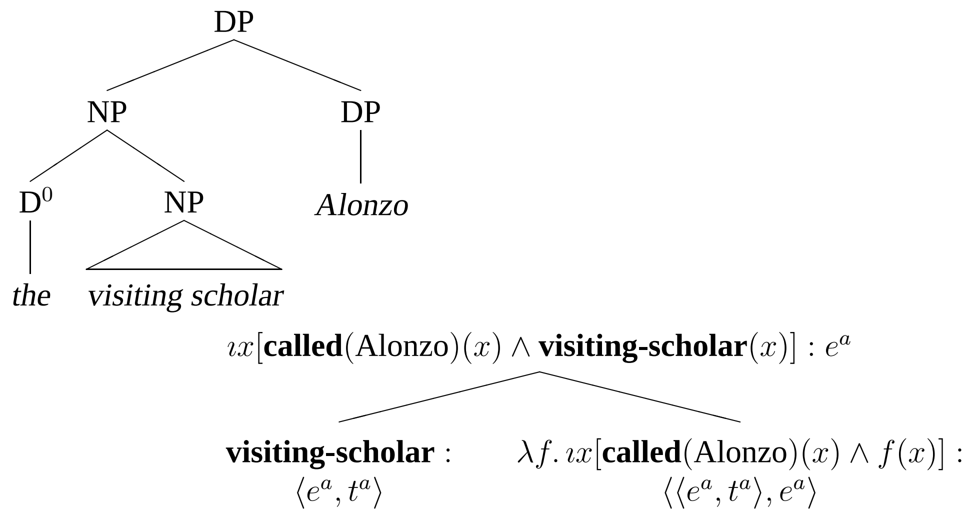

The analysis yields real dividends when we look at phenomena like extraction and variable binding. McCawley (1998:470) observes that it is impossible to extract from an NA, as seen in (3.107a–c). But extraction from some integrated appositives is fine, as in (3.107d).

(3.107) a. *[Which country]$_1$ did Tom put his shoes, a souvenir of a trip to $\mathbf{t}_1$, under the bed?
(McCawley 1998:470, (10a))

b. *[Italy]$_1$, Tom put his shoes, a souvenir of a trip to $\mathbf{t}_1$, under the bed?
(McCawley 1998:470, (10a′))

c. $^{??}$[The Illinois State Lottery]$_1$, my brother knows Albert Swenson, a recent winner of $\mathbf{t}_1$.
(McCawley 1998:470, (10b′))

d. $^{?}$[The balanced budget]$_1$, the president denounced {us/you} supporters of $\mathbf{t}_1$
(McCawley 1998:470, (11b))

Relative clauses provide an even clearer contrast. Integrated relatives do allow some extraction (the complex NP island constraint is not absolute), whereas supplementary relatives (as well as *As*-clauses; Potts 2002b and references therein) forbid all extraction to points external to them.

(3.108) a. What$_1$ the police arrested everyone who saw $\mathbf{t}_1$ was a video. (Postal 1998:9, (24a))

b. [That video]$_1$ the police arrested everyone who saw $\mathbf{t}_1$. (Postal 1998:9, (24b))

c. *What$_1$ the police arrested Eddie, who saw $\mathbf{t}_1$, was a video.

d. *[That video]$_1$ the police arrested Eddie, who saw $\mathbf{t}_1$.

NAs and supplementary relatives are not the only constructions that are distinguished largely by their intonational properties. VP *As*-parentheticals are the same way; the manner-adverbial counterparts, when accompanied by ellipsis, are distinguished from the parentheticals only by intonation:

(3.109) a. Alonzo built the canoe, as the foreman said he would (though not in the manner they expected him to build it).

b. Alonzo built the canoe as the foreman said he would (with an ax and elbow grease).

When we look at nonelided versions, the differences are sometimes more apparent:

(3.110) a. *Alonzo build the canoe, as the foreman said he would build the bicycle.

b. Alonzo build the canoe {as/in the way that} the foreman said he would build the bicycle.

Further probing reveals other differences. For instance, only VP *As*-parentheticals permit the kind of inversion discussed in Merchant 2002 and Potts 2002b:639.

(3.111) a. Alonzo built a canoe, as did Joan.

b. *Alonzo built a canoe as did Joan — with a hatchet.

c. Alonzo built a canoe as Joan did — with a hatchet.

Ideally, we would make all these differences stem from the presence or absence of the COMMA feature in the structures. I do not see a way to capture the inversion facts (Merchant's (2002) analysis of the inversion in (3.111a) is grounded in independent features of the syntax and holds also for comparatives, where there is no comma intonation). But we can ensure that the manner adverbials do not have a comma intonation associated with them. The types again provide the basis for the analysis; assume that manner-*As* works as in (3.112).

(3.112) $\textit{as}_{\textit{manner}} \rightsquigarrow \lambda p \lambda f \lambda x. \exists P[P(p) \wedge P(f(x))] : \langle t^a, \langle\langle e^a, t^a\rangle, \langle e^a, t^a\rangle\rangle\rangle$

It is worth looking at an illustration; let $\mathbf{as_m}$ abbreviate the lambda term in (3.112).

(3.113) a. Alonzo ran as Sue paddled.

b.

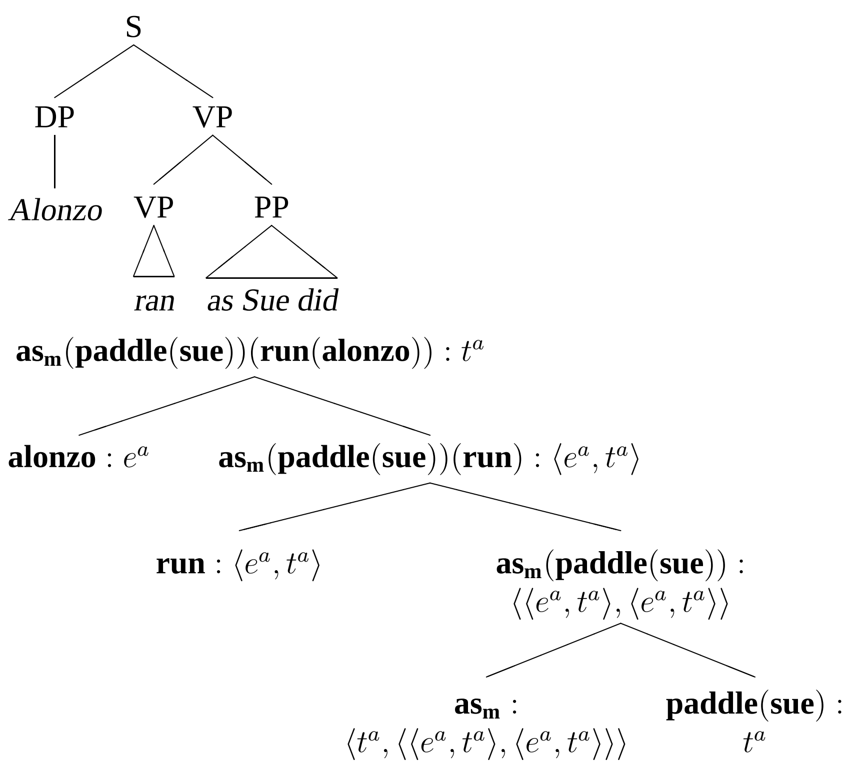

We ensure that these manner adverbials are not conventionally implicated by permitting no meaning for COMMA that takes terms with the type $\langle\langle e^a, t^a\rangle, \langle e^a, t^a\rangle\rangle$ to their CI correspondents. Stepping back, we see that the most concise way to define the entry for COMMA is as follows:

(3.114) COMMA $\rightsquigarrow \lambda X \lambda x.\, X(x) : \langle\langle \sigma^a, t^a\rangle, \langle \sigma^a, t^c\rangle\rangle$, where $\sigma \in \{e, t\}$

Because the manner-*as* clause has $\langle e^a, t^a\rangle$ outputs, no COMMA meaning is available for it. Its content cannot enter the CI dimension directly.

3.6.4 There are no inverted cases

It seems at first as though examples like (3.115a) are best analyzed as inverted variants of (3.115b).

(3.115) a. A former linguist, Edward Witten, is now the top-dog in string theory.
b. Edward Witten, a former linguist, is now the top-dog in string theory.

On this inverted analysis, we would have the following structure for (3.115a).

(3.116)

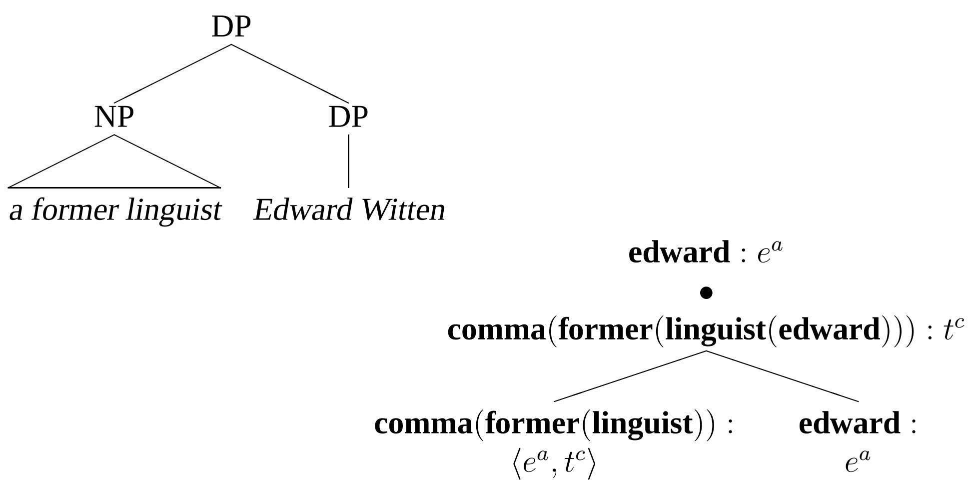

On this analysis, the parsetree for *A former linguist, Edward Witten* and *Edward Witten, a former linguist* are identical, since the parsetrees do not contain a linear-ordering relation. The differences are syntactic.

But all the evidence known to me suggests that this is incorrect. We must ensure that NAs always involve right-adjunction of the appositive to the anchor in the syntax, and moreover that this right-adjoined item is always the functor. Clear evidence for this analysis derives from existential constructions. Consider, first, the existential-*there* environment in (3.117).

(3.117) a. There was a former linguist at the party.
b. There was a former linguist, Ed Witten, at the party.
c. # There was Ed Witten at the party.
d. # There was Ed Witten, a former linguist, at the party.

We see definiteness effects only when the leftmost nominal is definite. These facts suggest that the anchor — the at-issue meaning contributor — is always on the left.

Additional support for this move derives from another existential construction, the one determined by *have* (Partee 1999). This environment imposes somewhat different restrictions than existential-*there*, but the prohibition on definite expressions is constant across both constructions. An inverted analysis of the complements to *have* in (3.118) wrongly predicts that (3.118d) is grammatical, because it is identical to (3.118b) on this approach.

(3.118) a. Ray had a student.
b. Ray had a student, Ed Witten.
c. # Ray had Ed Witten.
d. # Ray had Ed Witten, a student.

Finally, I offer some support from a slightly different domain: the two kinds of NA behave differently with regard to definiteness effects:

(3.119) a. Lance Armstrong, the cyclist, is from Texas, where he often trains with fellow cyclists.

b. # The cyclist, Lance Armstrong, is from Texas, where he often trains with fellow cyclists.

We have seen already that the definite inside an appositive need not be interpreted in its usual fashion (section 3.5.6). If *the cyclist* were the appositive (CI functor) in (3.119a) then we would expect this same shift in the meaning of the definite article to be possible. But it is not.

3.6.5 NAs summed up

Though more could be said, this draws to a close my study of NAs. It is worth noting that there seem to be substantive links between these constructions and the expressives of chapter 4. Aoun et al. (2001) claim that epithets in Lebanese Arabic take the form of NAs. From the perspective of the $\mathcal{L}_{\mathrm{CI}}$ description, this connection is unsurprising: expressives receive much the same treatment as supplements in the CI logic. In English, expressives generally lack the feature COMMA. The Lebanese Arabic facts indicate that this alone does not distinguish expressives from supplements. I pick up this discussion in much greater details in chapter 4. I close the present chapter by addressing a few supplemental expressions that meet the conditions for a CI analysis. These too provide a link with expressives, further suggesting that Grice's (1975) definition picks out a unified class of expressions.

3.7 Supplementary adverbs

The above analysis extends quite readily to a wide range of supplemental adverbs. In the main, this section addresses three classes of such expressions: speaker-oriented adverbs, topic-oriented adverbs, and utterance-modifiers. Jackendoff (1972) offered the first systematic study of these (and other) adverbs. Bellert's (1977) follow-up to Jackendoff's description is also important, as is the recent work by Nilsen (2001, 2003).

It is not my intention to attempt to improve on the syntactic descriptions of these adverbs that Jackendoff (1972) and later Huddleston and Pullum (2002) developed, nor do I aim for a fine-grained semantics for individual members of these classes. Rather, I want to show that $\mathcal{L}_{\mathrm{CI}}$ is a useful tool for getting at the well-known multidimensionality of the meanings that these expressions give rise to. The essential insight is due to Jackendoff (1972) and Bellert (1977). Of adverbs like *amazingly*, Bellert writes:

(3.120) "the adverb does not make part of just one proposition; in addition, we have a second proposition whose predicate (the adverb) evaluates the fact, event, or state of affairs denoted by S (sentence without the adverb). Sentences with evaluative adverbs express two asserted propositions [...]."
(Bellert 1977:342)

It seems to me that the missing ingredient in this description is the connection with Grice's (1975) definition of CIs. Below, I address various options for formalizing the claim in (3.120), concluding that the CI logic is the best of the bunch.

I also address the special restrictions on utterance-modifying adverbs, suggesting that they are best captured with a model-theoretic distinction.

3.7.1 Speaker-oriented adverbs

Consider, to begin, the speaker-oriented adverb *luckily* as used in (3.121).

(3.121) a. Luckily, Willie won the pool tournament.
b. Willie, luckily, won the pool tournament.
c. Willie won the pool tournament, luckily.

As usual, the commas should be taken seriously as intonational-phrase boundary markers. In each case, we have supplemental material on our hands. The adverb contributes a CI proposition, namely, the proposition that its propositional argument is positive. On this reading, Willie might have won in virtue of skill alone.

The examples contrast both intonationally and model-theoretically with the integrated adverbs in (3.122).

(3.122) a. Willie luckily won the pool tournament.
b. Willie won the pool tournament luckily.

In these examples, the adverb acts as a verb-phrase modifier. These examples assert that Willie won not by skill alone, but at least in part by luck. That is, *luckily* operates here as a manner adverbial. The meanings are not multidimensional.

These adverbs tend to have negated counterparts. If we use the negative counterpart of *luckily*, it becomes easy to see the two meanings working separately:

(3.123) Unluckily, Willie luckily won the pool tournament. (By a fluke, he sunk the winning 9-ball on the break; as a result, I had to pay $500 to Fats.)

So we again have excellent reason to focus attention on the comma intonation. It appears to be the main factor in differentiating the possible uses of *luckily*.

For manner-adverbial readings like (3.122), the semantics is straightforward. The adverb seems to have two possible extensional realizations: as a function from properties to properties and as a function from truth values to truth values; see (3.124).

(3.124) $\textit{luckily} \rightsquigarrow \begin{cases} \lambda f \lambda x.\, \textbf{lucky}(f(x)) : \langle\langle e^a, t^a\rangle, \langle e^a, t^a\rangle\rangle \\ \lambda p.\, \textbf{lucky}(p) : \langle t^a, t^a\rangle \end{cases}$

Both denotations play a central role in the analysis of the CI realizations of these adverbs. It seems that only the meaning that takes propositions into propositions has a CI semantics. This is exactly what the general meaning for COMMA in (3.114) predicts. The relevant permitted instantiation is (3.125), which takes us from an at-issue propositional modifier to a CI-based one.

(3.125) COMMA $\rightsquigarrow \lambda P.\, P : \langle\langle t^a, t^a\rangle, \langle t^a, t^c\rangle\rangle$

Here are two simple examples, the first involving regular manner-adverbial modification, the second involving a supplement with a CI semantics. Each begins with the same lexical meaning for *luckily*.

(3.126) Willie won the tournament luckily.

$$\textbf{lucky}(\textbf{win}(\textbf{the}(\textbf{tournament}))(\textbf{willie})) : t^a$$

$$\textbf{lucky} : \langle t^a, t^a\rangle \qquad \textbf{win}(\textbf{the}(\textbf{tournament}))(\textbf{willie}) : t^a$$

(3.127) Luckily, Willie won the tournament.

$$\mathbf{win}(\mathbf{the}(\mathbf{tournament}))(\mathbf{willie}) : t^a$$
$$\bullet$$
$$\mathbf{comma}(\mathbf{lucky})(\mathbf{win}(\mathbf{the}(\mathbf{tournament}))(\mathbf{willie})) : t^c$$

$$\mathbf{comma}(\mathbf{lucky}) : \langle t^a, t^c \rangle \qquad \mathbf{win}(\mathbf{the\text{-}tournament})(\mathbf{willie}) : t^a$$

$$\mathbf{lucky} : \langle t^a, t^a \rangle$$

From the point of view of the logic, speaker-oriented adverbs justify the (slightly) general meaning for COMMA in (3.114).

We've now seen how the $\mathcal{L}_{\mathrm{CI}}$ description works. It is worth taking time to explore other options. To do this, I offer first a naturally-occurring example, one in which the speaker-orientation of the adverb is especially clear:

(3.128) "After first agreeing to lend me a modem to test, Motorola changed its mind and said that, amazingly, it had none to spare".[7]

The context indicates that Motorola did not characterize its claimed modem shortage with *amazingly*. The speaker used the adverb to express skepticism. The content of Motorola's utterance is simply that it had no modems to spare.

As we saw already in the quotation from Bellert 1977 above, the multidimensionality of these meanings has already been acknowledged. But neither Jackendoff nor Bellert makes the connection with CIs. Rather, their descriptions seem to suggest one of two analyses: (*i*) the adverb brings with it a factive presupposition that gives rise to the sense that more than one proposition is expressed; or (*ii*) the adverb is a function

[7] <http://www.hamline.edu/apakabar/basisdata/1997/03/21/0066.html>.

that takes propositions to product-types. We can, in fact, dispatch with both alternatives.

First, consider what a presuppositional denotation would look like:

(3.129) *amazingly* $\rightsquigarrow \lambda \mathbf{p} \lambda w \ddagger \mathbf{p}(w) = 1 \ddagger . \textbf{amazing}_w(\mathbf{p})$
(*denoting a partial function from propositions into propositions*)

This names a partial function from propositions into propositions; the expression inside the doubled daggers specifies the domain condition. At a technical level, the proposal closely resembles the approach taken in Potts 2002a,b. However, various arguments suggest that it is incorrect. A compelling and straightforward one derives from pairs like the following, which indicate that we do need meanings like the one in (3.129).

(3.130) a. Motorola said that it is amazing that it has no spare modems.
b. Motorola said that, amazingly, it has no spare modems.

Example (3.130a) commits Motorola to the characterization of its modem shortage as amazing. But (3.130b) does not. The speaker characterizes this shortage with *amazingly* in such a way as to conversationally implicate skepticism. To achieve this in presuppositional terms, we would have to stipulate that the presupposition is satisfied in the actual world (as in Potts 2002a,b).

There is a second alternative to the $\mathcal{L}_{\text{CI}}$ description, one that seems to match the descriptions of Bellert and Jackendoff more closely. On this view, we make *amazingly* denote a function taking propositions into pairs of propositions:

(3.131) *amazingly* $\rightsquigarrow$
$$\lambda \mathbf{p}. \left\langle \begin{array}{c} \mathbf{p}, \\ \lambda w. \textbf{amazing}_w(\mathbf{p}) \end{array} \right\rangle : \langle \langle s^a, t^a \rangle, \langle \langle s^a, t^a \rangle \times \langle s^a, t^a \rangle \rangle \rangle$$

This denotation has one important advantage over the presupposition-trigger treatment in (3.129): it does not require that the input proposition be entailed by the common ground. It rightly allows that a clause containing *amazingly* might introduce entirely new information.

However, it seems that this treatment does not get the scope facts right, in that it wrongly embeds the contribution of the adverb inside certain intensional contexts. For instance, we saw in chapter 1 that *say* must be allowed to take tuples of propositions into tuples of propositions, in order to capture Bach's (1999) observations that the secondary proposition contributed by *but* (and synonyms) must often become part of the argument to this verb. I refer to chapter 6 for a fuller discussion of phenomena resembling those that Bach addresses. For now, suffice it to say that those assumptions allow for derivations like the following (recall that $\overrightarrow{\langle s^a, t^a \rangle}_2$ abbreviates $\langle \langle s^a, t^a \rangle \times \langle s^a, t^a \rangle \rangle$):

(3.132)

$$\textbf{say}\left\langle \begin{array}{c}\textbf{modemless}(\textbf{motorola}),\\ \textbf{amazingly}(\textbf{modemless}(\textbf{motorola}))\end{array}\right\rangle : \langle e^a, \overrightarrow{\langle s^a, t^a \rangle}_2 \rangle$$

$$\textbf{say} : \langle \overrightarrow{\langle s^a, t^a \rangle}_2, \langle e^a, \overrightarrow{\langle s^a, t^a \rangle}_2 \rangle \rangle \qquad \left\langle \begin{array}{c}\textbf{modemless}(\textbf{motorola}),\\ \textbf{amazingly}(\textbf{modemless}(\textbf{motorola}))\end{array}\right\rangle : \overrightarrow{\langle s^a, t^a \rangle}_2$$

$$\textbf{amazingly} : \langle \langle s^a, t^a \rangle, \overrightarrow{\langle s^a, t^a \rangle}_2 \rangle \qquad \textbf{modemless}(\textbf{motorola}) : \langle s^a, t^a \rangle$$

Like the presuppositional treatment, this analysis wrongly predicts that Motorola characterized its modem shortage as amazing. This is, as noted above, quite clearly not the intended reading of the example on which this parsetree is based. The main virtue of the CI analysis in this area is that it neatly separates out the contribution of the adverb, giving it root-level force.

3.7.2 Topic-oriented adverbs

My analysis of speaker-oriented adverbs is easily and I think fruitfully extended to Jackendoff's 'subject-oriented' adverbs:

(3.133) a. Thoughtfully, Edna washed the dishes.
b. Edna, thoughtfully, washed the dishes.
c. Edna washed the dishes, thoughtfully.

The only new challenge posed by this class of adverbs is that we appear to need access to the subject, as it seems to be one of the semantic arguments to the adverb (Bellert 1977). There are two options for how to do this. The more complicated of the two would involve accessing the subject argument directly. This would be technically challenging, because it seems that the argument to the adverb is a proposition. There is no simple way to recover from this semantic object the denotation of the subject.

Fortunately, these machinations seem unjustified. Though there is a great deal of prescriptive pressure to have the argument to sentence-initial free-adjuncts of this kind be the subject, this is not in fact how speakers behave. Arnold Zwicky (p.c., 2/03) offers examples like those in (3.134a–b), which he culled from independent sources. I myself found (3.134c).

(3.134) a. [The narrator is describing his time in the Visalia County, CA, Jail.]
"After more than a month in jail, my mother posted bond, bless her soul."[8]

[8] Kipnis, Aaron. 1999. *Angry Young Men*. San Francisco: Jossey-Bass (p. 130).

b. "Signed by Columbia Records in 1999, his first album was never released."[9]

c. "Having decided on the Turing machine as the basic computing device for the theory, the idea is to measure the efficiency of an algorithm by the number of steps (i.e., Turing-machine steps) taken to complete the calculation."[10]

In each case, the initial adjunct is syntactically incomplete, but it in fact has a propositional denotation. Stump (1985) observes that the relationship between these phrases and the clauses they adjoin to is extremely variable.

The requirements on the argument to 'subject-oriented' adverbs seems the same as it is for free adjuncts. Consider the parallels between (3.134) and (3.135).

(3.135) a. "Physically, the keyboard is smaller than I expected, and extremely well built — there's no creaking or flexing. The keys look as if they will last well — including their paint. Thoughtfully, there is a clip-on cover for the connector while not in use."[11]

b. "The music, while well constructed, is rather annoying after a while, with a lack of any instantly recognizable tunes apparent. But, thoughtfully, there is an option to turn the sound off at any time during the game, so the rather twee sound effects and jauntily repetitive soundtrack won't annoy the parents too much when they've passed out on the settee from too much Christmas pud."[12]

c. "What is the function of the marking in the highest clause? Tentatively, it signals the left edge of a nominalized relative clause-type syntactic constituent."[13]

---

[9] *Fresh Air*. NPR. February 21, 2003. Interview with the rap star Fifty Cent.

[10] Devlin, Keith. 1998. *Mathematics: The New Golden Age*. New and revised edition. New York: Columbia University Press (p. 292).

[11] <http://www.pdatweaks.com/reviews.php?itemid=238>

[12] <http://www.gamesdomain.com/gameboy/previews/Santa_Claus_Junior.html>

[13] Gerassimova, Veronica. 2002. Handout for dissertation research report, Stanford, 11/26/02. My thanks to Arnold Zwicky for finding this one.

Thus, it seems that we seek a looser connection between the 'subject-oriented' adverb and the subject itself, as in (3.136).

(3.136) a. Thoughtfully, the batteries were included.

b.

$$\textbf{included}(\textbf{the}(\textbf{batteries}))(x_1) : \langle s^a, t^a \rangle$$

$$\bullet$$

$$\textbf{thoughtfully}(\textbf{included}(\textbf{the}(\textbf{batteries}))(x_1))(x_1) : \langle s^a, t^c \rangle$$

$$\textbf{comma}(\lambda \mathbf{p}.\, \textbf{thoughtfully}(\mathbf{p})(x_1)) : \langle \langle s^a, t^a \rangle, \langle s^a, t^c \rangle \rangle \qquad \textbf{included}(\textbf{the}(\textbf{batteries}))(x_1) : \langle s^a, t^a \rangle$$

$$\lambda \mathbf{p}.\, \textbf{thoughtfully}(\mathbf{p})(x_1) : \langle \langle s^a, t^a \rangle, \langle s^a, t^a \rangle \rangle$$

Here, the most salient reading seems to be the one on which x_1 is both the thoughtful individual and the individual who included the batteries. The denotation of x_1 is a discourse topic. For this reason, I adopt the coverterm *topic-oriented* for this class of adverbs, expanding the insight of Arnold Zwicky (p.c., 1/03) that the meanings of preposed modifiers like those in (3.134) are generally saturated by a discourse topic. For both classes of construction, the tendency to think that the argument to the adverb must be the subject no doubt derives from the fact that, in English, topics tend to be expressed in subject position, and subjects tend to denote topics.

3.7.3 Utterance-modifiers

The most intriguing supplemental adverbials are the *utterance-modifying* ('pragmatic', 'second-order') adverbs, exemplified in (3.137).

(3.137) a. Frankly, I am sick of your complaining.
b. Between you and me, Ed wears a toupee.
c. Soldier to soldier, I find army-issue underwear uncomfortable.
d. In case you're interested, Ed fled.

These modifiers are discussed in Thorne 1972, Jackendoff 1972, Bellert 1977:349, and Bach 1999:§5. They tend to appear sentence-initially, though sentence-final occurrences are also found:

(3.138) a. I am sick of your complaining, frankly.
b. Ed fled, in case you're interested.

However, genuine embedding is not possible. This is evident in two ways. First, when we place these adverbs inside even matrix verb phrases, they change their meaning slightly. Consider, for instance, this pair:

(3.139) a. Confidentially, Ed mentioned the bribe.
b. Ed confidentially mentioned the bribe.

For (3.139b), Ed must have mentioned the bribe in confidence. The speaker, however, makes no request for confidentiality. One could follow an utterance of (3.139b) with something like "But I don't know what he's being so secretive about". In contrast, *confidentially* in (3.139a) works to impart the speaker's directive that his addressee keep quiet the fact that Ed mentioned the bribe. It might be that Ed made no such request — he might have been blabbing to everyone about the sordid deal in question. (It is possible that a (3.139b)-type reading is available for (3.139a). If so, it is not prominent.)

This is not the only sense in which embedding changes the possible readings. A similar shift is enforced when we move to complement clauses; in some cases, embedding leads to outright ungrammaticality; see (3.140), from Bach 1999:358, (32)–(34).

(3.140) a. Confidentially, Al's wife is having an affair.

b. # Bill said that, confidentially, Al's wife is having an affair.

c. Man to man, your wife is having an affair.

d. # Bill said to Al that, man to man, his wife was having an affair.

e. * Bill said to Al that, in case he was interested, Al's wife was having an affair.

If (3.140b) and (3.140d) have any readings at all, they are ones in which the adverbials work as integrated modifiers characterizing the way in which Al's wife is having an affair. Since both readings are unlikely, the examples are infelicitous. The central datum to be accounted for is that the reading on which *confidentially* is a speaker-oriented contribution (an instruction to the hearer not to spread the word) is unavailable. This distinguishes these adverbs from the other supplements we have looked at so far, which are syntactically embeddable though semantically unembeddable. Here, even syntactic embedding is blocked.

At first, one might think that this merely calls for a syntactic constraint on where these adverbs can appear. It would be relatively easy to formulate the condition: we would simply state that the mother node of any utterance-modifying adverbial is the root node. We would not, on this approach, need to mention semantic parsetrees. However, the discussion of Thorne (1972) shows that this would not be a complete account. Thorne's main descriptive contribution is the observation that utterance modifiers can appear inside supplementary relatives. His observations extend to a wide range of clausal supplements, as seen in (3.141).

(3.141) a. Eddie, who, just between you and me, has been stealing from the collection plate, has just been promoted.

b. We interviewed Lance, quite frankly the best cyclist in the world right now, about his plans for the future.

c. That Ames was a spy, as the *Times* reported early on, in case you're interested, was a shock to the FBI.

These examples point up the inadequacy of the stipulation that the adverb must be the daughter to the root node. This criticism would hold even if we were to adopt the view (discussed in detail in chapter 5) that supplements are adjoined at the root node, because we have syntactic embedding in both supplementary relatives and *As*-parentheticals.

It seems to me that these examples indicate quite clearly that utterance-modifiers are restricted to matrix occurrences because they require arguments that have main clause force. On the current approach, supplements and main clauses are unified in this respect in virtue of the interpretive rule in (3.18), which I repeat here:

(3.142) **parsetree interpretation**

Let $\mathcal{T}$ be a semantic parsetree with the at-issue term $\alpha : \sigma^a$ on its root node, and distinct terms $\beta_1 : \langle s^a, t^c \rangle,\ \ldots,\ \beta_n : \langle s^a, t^c \rangle$ on nodes in it (extensionally, $\beta_1 : t^c,\ \ldots,\ \beta_n : t^c$). Then the interpretation of $\mathcal{T}$ is the tuple

$$\left\langle [\![\alpha : \sigma^a]\!]^{\mathcal{M}_i,g}, [\![\beta_1 : \langle s^a, t^c \rangle]\!]^{\mathcal{M}_i,g}, \ldots, [\![\beta_n : \langle s^a, t^c \rangle]\!]^{\mathcal{M}_i,g} \right\rangle$$

where $[\![\cdot]\!]^{\mathcal{M}_i,g}$ is the interpretation function, taking formulae of the meaning language to the interpreted structure $\mathcal{M}_i$, relative to a variable assignment g.

Though CI terms might be embedded inside the semantic parsetree, they are interpreted as though they had main clause force.

Bellert (1977:349) makes an important observation about utterance-modifying adverbials (which she calls 'pragmatic adverbs'): they can optionally occur with *speaking*. Thus, alongside *frankly* we have *frankly speaking* and *speaking frankly*; alongside *soldier to soldier* we have *speaking soldier to soldier*, and so forth. This is not possible for adverbs in most other classes:

(3.143)	a.	*possibly speaking	*speaking possibly
	b.	*obviously speaking	*speaking obviously
	c.	*amazingly speaking	*speaking amazingly

We can intuit at least part of the semantics for *speaking*: it has as one of its arguments the speaker of the clause.

I think that we should take seriously the intuition behind the labels 'utterance modifier', 'second-order', and 'pragmatic'. What we want to say is that an adverb like *frankly* modifies the relation between a speaker and a particular sentence. Bellert's observation that *speaking* can appear with these modifiers suggests that it is truly the utterance relation that we modify with utterance modifiers. To achieve this, I exploit the upper layer of the *discourse structures* defined in chapter 2. The definition of those structures is repeated here:

(3.144) A *discourse structure* is a tuple $\mathcal{D} = (A, \mathbf{D}, D_u, \mathfrak{M}, \hbar, V_{\mathcal{D}})$, where

a. $A = \{a_1, a_2, \ldots\}$ is a set of discourse participants.

b. $D_u = \{\mathcal{S}_1, \mathcal{S}_2, \ldots\}$ is a set of sentences, the domain of u. Each $\mathcal{S}$ is a pair $(\mathcal{T}^s, \mathcal{T}^m)$, in which $\mathcal{T}^s$ is a syntactic structure and $\mathcal{T}^m$ is its associated semantic parsetree (as defined in (2.36)). D_u contains a subset $D_d = \{\mathcal{D}_1, \mathcal{D}_n, \ldots\}$ of declaratives (the domain of d) and a subset $D_q = \{\mathcal{Q}_1, \mathcal{Q}_n, \ldots\}$ of interrogatives (the domain of q). $D_q \cap D_d = \emptyset$.

c. $\mathbf{D}$ is a set of domains, as defined in (2.61); $A \subseteq D_e$.

d. $\mathfrak{M} = \{\mathcal{M}_1, \mathcal{M}_2, \ldots\}$ is a set of intensional models, as defined in (A.21). All $\mathcal{M}_i \in \mathfrak{M}$ have $\mathbf{D}$ as their set of domains.

e. $\hbar$ is a function that takes each $a_i \in A$ to the model $\mathcal{M}_i \in \mathfrak{M}$, where $\mathcal{M}_i$ can be viewed as the world-view of a_i.

f. $V_{\mathcal{D}}$ is a valuation function, taking constants of $\mathcal{L}_U$ to functions formed from objects in $D_e \cup D_u \cup \{0, 1\}$, constrained so that if α is of type σ, then $V_{\mathcal{D}}(\alpha) \in D_\sigma$.

In brief, these structures contain a set A of discourse participants, each associated with his own intensional model in $\mathfrak{M}$, where each of those models shares a single set of domains $\mathbf{D}$ (but possibly views that domain differently). These structures also contain a set D_u of sentences, which is broken up into a set of declaratives and a set of interrogatives. Each object in D_u is a pair $\mathcal{S} = (\mathcal{T}^s, \mathcal{T}^m)$, where $\mathcal{T}^s$ is a syntactic structures and $\mathcal{T}^m$ is a semantic parsetree, as defined in chapter 2, section 2.6 (page 80). These parsetrees are labelled with terms of $\mathcal{L}_{\text{CI}}$, according to the tree-admissibility conditions of the logic.

The logic $\mathcal{L}_U$ permits us to talk about the upper layer of these structures. In particular, we can defined relations between members of A and members of D_u. The term $\ulcorner$**utter**$\urcorner$ is an important term of $\mathcal{L}_U$. It is defined as in (3.145), in which S serves as a variable over objects of type u, where u is the type of objects in D_u.

(3.145) *speaking* $\rightsquigarrow \ulcorner\textbf{utter}\urcorner \stackrel{def}{=} \lambda S \lambda x. \ulcorner\textbf{utter}\urcorner(S)(x) : \langle u, \langle e, t\rangle\rangle$

As in chapter 2, I relativize the interpretation function $[\![\cdot]\!]^{\mathcal{D},s,a}$ for a discourse structure $\mathcal{D}$ to a speaker s and an addressee a, where $s, a \in A$:

(3.146) a. $[\![\varphi]\!]^{\mathcal{D},s,a} = V_{\mathcal{D}}(\varphi)$ if φ is a formula of $\mathcal{L}_U$.

b. $[\![\mathcal{S}]\!]^{\mathcal{D},s,a}$ = the value of $\mathcal{S}$ determined by (3.18) if $\mathcal{S}$ is a parsetree for $\mathcal{L}_{\text{CI}}$.

We now have the pieces in place to offer denotations for utterance modifiers. In (3.147), I offer an analysis of *frankly*.

(3.147) a. $[\![\ulcorner\textbf{the-speaker}\urcorner]\!]^{\mathcal{D},s,a} = s$

b. $\ulcorner\textbf{frankly}\urcorner : \langle\langle d, \langle e, t\rangle\rangle, \langle d, t\rangle\rangle$

c. $\lambda S. \ulcorner\textbf{frankly}\urcorner(\ulcorner\textbf{utter}\urcorner(S))(\ulcorner\textbf{the-speaker}\urcorner) : \langle d, t\rangle$

$\lambda \mathcal{U} \lambda S. \ulcorner\textbf{frankly}\urcorner(\mathcal{U}(S))(\ulcorner\textbf{the-speaker}\urcorner) :$
$\langle\langle u, \langle e, t\rangle\rangle, \langle d, t\rangle\rangle$

$\ulcorner\textbf{utter}\urcorner :$
$\langle u, \langle e, t\rangle\rangle$

The only surprising thing about this analysis is that the mother node denotes a function that requires entire parsetrees as its argument. Thus, it is not a member of the set of parsetrees defined in (2.36). A perspicuous way of representing the link between (3.147) and those structures is as in (3.148). The dashed line represents the action of the interpretation function on ⌜Ed fled⌝, which is a constant of $\mathcal{L}_U$ that denotes the member of D_u represented in (3.148b).

(3.148) a.

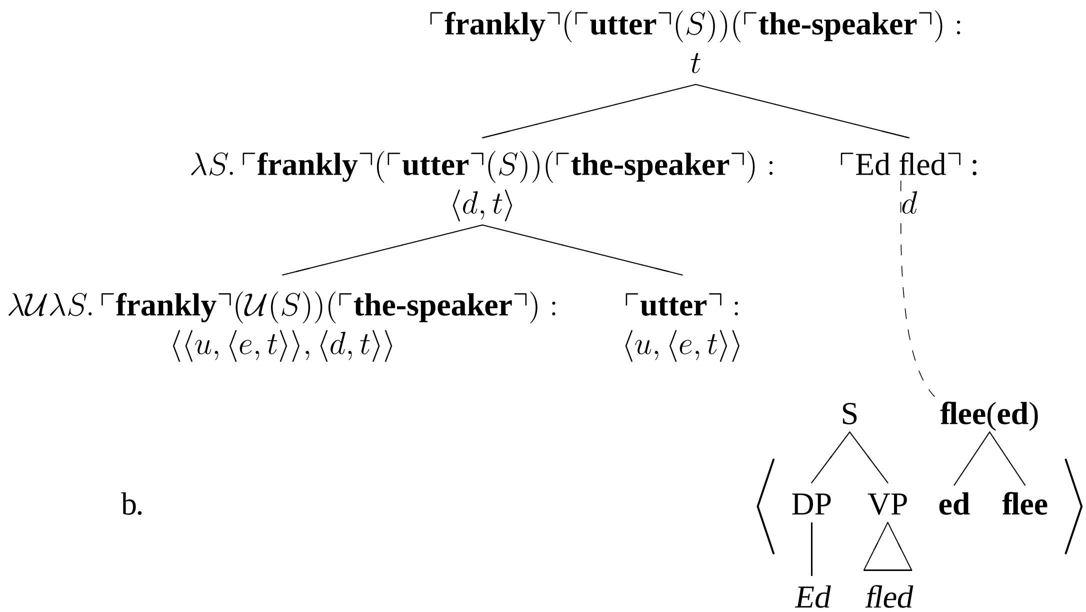

It might seem surprising that the adverb *frankly* is not part of the sentence it modifies in either the syntactic or the semantic sense. But this move is necessary: it is the most direct possibly formalization of the correct intuition that these items place conditions on (i.e., modify) the utterance relation.

The analysis easily accounts for the unembeddability of these operators that is evident in (3.140). It is untrue that, in example (3.140b), the speaker is in the utterance relation with *Al's wife is having an affair.* This is an utterance that he attributes to Bill, not one that he makes himself. Since utterance modifiers like this are defined as relations between the speaker and a sentence, we cannot use them to form relations between entities other than the speaker. For (3.140b), this means that we cannot relativize the content of *confidentially* to Bill.

This general approach finds some support in examples like (3.149), all of which have first-person subjects.

(3.149) a. I'd say that, just (speaking) between you and me, Ed fled with the winnings.

b. I feel that, quite frankly (speaking), Ed is not trustworthy.

c. I swear that, (speaking) man to man, I did not sell your chihuahua into slavery.

The improvement makes sense when we consider that, if someone says *I'd say that S*, we invariably regard him as having said *S*. Thus, the following discourse is infelicitous:

(3.150) I'd say that Ed fled with the winnings. $^{\#}$But I did not (do not) say that Ed fled with the winnings.

Technically speaking, this means that we should enforce the following condition:

(3.151) For all x and $\ulcorner$S$\urcorner$,
if $[\![\ulcorner\textbf{utter}\urcorner(\ulcorner x \text{ says S}\urcorner)(x)]\!]^{\mathcal{D},s,a} = 1$, then $[\![\ulcorner\textbf{utter}\urcorner(\ulcorner\text{S}\urcorner)(x)]\!]^{\mathcal{D},s,a} = 1$.

Similar meaning postulates are necessary to ensure that speakers are seen has having uttered their supplements, so that we can accommodate Thorne's (1972) observation that utterance modifiers can appear inside them (see (3.141) above).

Utterance modifiers also point up the need for dividing the set D_u into a set D_q of interrogatives and a set D_d of declaratives. Bill Ladusaw observes (p.c., 5/03) that when an utterance modifier takes an interrogative argument, the meaning changes, becoming, in a sense, addressee-oriented. I repeat the following examples from chapter 1 (example (1.18)).

(3.152) a. Confidentially, is Al having an affair?
$\approx$I promise to keep the answer to *Is Al having an affair?* a secret.

b. Honestly, has Ed fled?
$\approx$Provide me with an honest answer to the question *Has Ed fled?*

To describe these facts, we need only assume that utterance modifiers are ambiguous. We have, for instance, the following pair of meanings for *honestly speaking*.

(3.153) *honestly speaking* $\rightsquigarrow$

$$\begin{cases} \lambda S.\ulcorner\textbf{honestly}\urcorner(\ulcorner\textbf{utter}\urcorner(S))(\textbf{the-speaker}) : \\ \hfill \langle\langle u, \langle e,t\rangle\rangle, \langle d,t\rangle\rangle \\ \lambda Q.\ulcorner\textbf{honestly}\urcorner(\ulcorner\textbf{utter}\urcorner(\ulcorner\textbf{answer}\urcorner(Q)))(\textbf{the-addressee}) : \\ \hfill \langle\langle u, \langle e,t\rangle\rangle, \langle q,t\rangle\rangle \end{cases}$$

Recall that q for the type of members of D_q. The denotation of **answer** applies to an interrogative Q to return a felicitous answer to Q. Thus, the second meaning is true only if the addressee enters into the honest-utterance relation with the answer to the question Q. The first meaning is identical in relevant respects to the meaning for *frankly* defined in (3.147b).

We could expand this account considerably. Discourse structures pave the way for a semantics for a host of words and phrases that speakers use to talk about and organize discourses. Phrases like *one the one hand*, *in contrast*, and *first*, when used as devices for organizing a series of sentences, can be modelled using these structures and the layered logic afforded by the combination of $\mathcal{L}_{\text{CI}}$ and $\mathcal{L}_U$ that is created by the interpretation function defined in (3.146). These extensions and others are the subject of Potts 2003b.

3.8 Conclusion

The main achievement of this chapter is the substantive links it develops between Grice's (1975) definition of CIs and the world of supplements. Much of the descriptive work is given over to NAs, but mainly because these nicely illustrate the way that a logic like $\mathcal{L}_{\text{CI}}$ is useful for describing the semantics of supplements while at the same time respecting the wealth of evidence for a conservative, modifier-adjunction-based theory of the syntax.

The account also fills a noticeable lacuna in the treatments of Potts 2002a,b in taking seriously the idea that comma intonation is a central part of what makes supplements special. In those earlier papers, comma intonation is essentially regarded as a mere reflex of the special semantics of these expressions. In the present account, it takes center stage in the sense that **comma**, the translation of the feature COMMA, takes regular content and returns CI content, thereby employing the CI dimension of $\mathcal{L}_{\text{CI}}$ in a substantive way.

Chapter 4 is devoted to a preliminary treatment of expressive meanings as CI contributors. The connection with the present chapter might seem indirect at first, since expressives tend to lack any semblance of supplementation. But recent work by Aoun et al. (2001) suggests that the links might in fact might be subtle and important in their own right. For instance, it seems that some, perhaps all, epithets in Lebanese Arabic have the form of NAs. What's more, there is, intuitively, a link between utterance modifiers and expressive meaning; it is possible that interjections like *damn* in *Damn, I broke the toaster* are a species of utterance-modifier.

In chapter 5, I pause to consider a more radical syntax for supplements, one that could be regarded as an alternative to the CI hypothesis. I reject the notion that it is useful for the constructions addressed here, but it might be that it is suitable for certain more peripheral expressions (*of course*, *for example*). Hence, the discussion has a positive aim.

Chapter 4

Expressive content

4.1 Composition and denotation

In the area of *expressives*, exemplified briefly in (4.1), we find excellent additional support for conventional implicatures (CIs).

(4.1) a. **Japanese verbal (subject) honorification**

Sensei-wa eigo ga <u>o-wakari-ni</u> <u>nar-u</u>
the-teacher-SUBJ *English* NOM HON-*understanding*-DAT *become*-IMP
'The teacher understands English.' (Toribio 1990:535, (1a))

b. **expressive attributive adjectives**

Shut that <u>blasted</u> window! (Cruse 1986:272,(19b))

c. **epithets**

saami ha-l-maʒduub nəse l-mawʕad
Sami 3-*the-idiot*.SM *forgot*.3SM *the-appointment*
'Sami, this idiot, forgot the appointment.' (Aoun et al. 2001:385, (37a))

The content of these expressions and a range of others reviewed below is speaker-oriented and intuitively independent from the at-issue content (Soames 2002:57–58).

Here, more than anywhere else, the idea that CI items comment upon an asserted core, providing a means for a bit of editorializing on the part of speakers, seems apt. As a result, the CI logic proves an excellent tool for managing this kind of content. The present chapter is largely devoted to substantiating this claim. The discussion mainly concerns the semantic combinatorics. The analyses are interpretable in the discourse structures of chapter 2, section 2.8. But a great many interpretations are possible.

A methodological point lies behind this observation: managing the content is separable from specifying what the content is. We can do a great deal of semantic investigation without specifying precisely what the models look like, simply because a large and important part of semantic theorizing is essentially syntactic (or combinatoric), addressing issues of modes of composition. This work can be done without a firm structural definition in place. This is in fact a familiar point: a precedent is the Amsterdam-style reinterpretation of first-order logical formulae in a dynamic fashion, which moves the denotation domain from sets of assignments (as in static semantics) to sets of input–output assignment pairs (the dynamic idea). More radically, all students of the direct interpretation mode have seen the models change over the years while the natural language stayed (for all intents and purposes) constant. We needn't fiddle with anything about English syntax to move from a Montagovian model based in total interpretation functions to the partial models of situation semantics. Similarly, we can add types for places, events, vectors, kinds, and the like without adjusting the fundamental assumptions about how composition proceeds.

This chapter divides into two main sections. The first, section 4.2, gathers together various insights that linguists have offered about expressives. The aim is a working definition of *expressive content*. Remarkably, the definition we arrive at matches Grice's (1975) definition of CIs, justifying an extension of the CI logic into this new domain.

I make that extension in section 4.3. Concentrating on the content of expressive attributive adjectives like *damn* and the descriptive part of epithets, I show how the CI logic can manage expressive content. In tandem with the system they are embedded in, the proposed lexical entries deliver a range of correct results about the limitations on these items. Facts involving epithets with at-issue quantifier associations suggest the addition of the insights of Gawron's (1996) restriction logic to this new domain.

This investigation is followed by a look at Japanese honorifics and the German discourse subjunctive, Konjunktiv I.

Along the way, there is much work to be done. The CI hypothesis entails the view that expressive content is always speaker-oriented. Apparent counterexamples to this claim are arguably not actual counterexamples. Their intonational contour and discourse-felicity conditions suggest that they are quotative. This perspective on matters reveals such readings to involve, not narrow or intermediate scope for the expressive content, but rather a total lack of scopal relations at all, engendered in this case by the essential referentiality of quotative utterances.

A motto for this chapter is that we lack compelling reasons to assume that *any* content falls outside the bounds of the modelling techniques available to natural-language semanticists. Thus, for example, I do not subscribe to the position advocated in (4.2).

(4.2) "The question as to what properties are associated with a person who is referred to through honorifics is not a problem to which a grammatical description is addressed, though it is an interesting matter for sociolinguistic researches." (Harada 1976:500)

Broadly speaking, Japanese argument-oriented honorifics signal that a certain entity mentioned in the sentence is socially superior to the speaker. I see no reason to assume that we cannot ask formally-informed questions about this relation — is it transitive? reflexive? — in the same way that we could ask questions about the meaning of a two-place relation like the one named by *socially superior to*. Rather than foisting these questions off onto sociolinguistics, people who write grammars should address them head-on. The result can be enlightening, and it can have significant consequences for the shape of semantic theory.

Alternative classifications of claimed CI content always loom large, especially those that seek to assimilate it to at-issue entailment. Chapter 5 addresses a syntactic alternative along these lines. That account, which is based in a new view of natural language syntactic tree geometry, is much less plausible in the area of expressives, which tend to be syntactically and intonationally integrated, and moreover fall into existing syntactic classes. But here we see the potential utility of a sophisticated scope-shifting account. Section 4.5 spells out the challenges facing an account in this vein, arguing that important aspects of Grice's (1975) definition cannot be captured in these terms without conceding all points of substance to the CI approach. The main utility of this discussion is not to discourage alternative formalizations, but rather to show that a few key concepts are bound to turn up in any adequate description.

4.2 A working definition

Part of the task of this early investigation is to come up with a working definition of *expressive*, thereby providing us with a means for deciding in a principled way what should and should not be judged relevant to the discussion. The best way to do this is to balance existing insights about expressive content with our intuitive sense for what sort of expression this term characterizes.

I repeat the definition of CIs in (4.3).

(4.3)
a. CIs are part of the conventional (lexical) meaning of words.
b. CIs are commitments, and thus give rise to entailments.
c. These commitments are made by *the speaker of the utterance* "by virtue of the meaning of" the words he chooses.
d. CIs are logically and compositionally independent of what is "*said* (in the favored sense)", i.e., independent of the at-issue entailments.

There are striking parallels between this definition and the observations that have been made in the literature about the nature of expressive content. Although the issue of whether expressives are best thought of as entailments is a tricky one, the other criteria are easily matched with claims about what makes expressive special. The following are drawn from discussions by Cruse (1986), Kaplan (1999), and Löbner (2002):

(4.4)
a. **lexicality**

"expressive meaning is part of the lexical meaning of certain expressions, a semantic quality of words and phrases"

(Löbner 2002:32)

b. **entailment**

"the aspects of meaning under discussion, in particular, the semantic information displayed by expressives, can have consequences for the notion of logical validity"

(Kaplan 1999:13)

c. **speaker-orientation**

"Another characteristic distinguishing expressive meaning from propositional meaning is that it is valid only for the utterer, at the time and place of utterance. This limitation it shares with, for instance, a smile, a frown, a gesture of impatience [. . .]."

(Cruse 1986:272)

"the prior discussion should make us cautious about always accepting as legitimate the demand of a report in indirect discourse"

(Kaplan 1999:8)

d. **independence**

"Expressive meaning carried by a lexical item in a statement plays no role in determining its truth-conditions."

(Cruse 1986:272)

The lexicality property suggests a narrowly grammatical treatment. If we view $\mathcal{L}_{\text{CI}}$ as a set of parsetrees whose terminals are labeled with lexical meanings, then we should be able to derive expressive meanings by simply including them in individual lexical items. The claimed utility of $\mathcal{L}_{\text{CI}}$ in this respect is of course greatly bolstered by the multidimensionality property (4.4d) and the speaker-orientation property (4.4c), as these are the properties that $\mathcal{L}_{\text{CI}}$ is designed to capture.

Do expressives count as entailments? The usual test for this is deniability — can a speaker use the item in question and then, in the same context, deny the content it expresses? For expressives, the answers is clearly negative; Cruse (1986:271) observes:

(4.5) "At this risk of being thought presumptuous, one could challenge the veracity of 15a [=*I just felt a sudden sharp pain*]; [. . .] it would make little sense to challenge 15b [=*Ouch*]."

We could say the same for a use of *damn* in *the damn Republicans.* This is not at all like *the gray-haired Republicans*, in that *gray-haired* could incorrectly characterize some set of entities, whereas *damn* does not offer this kind of speaker-independent meaning. A sincere utterance of *damn* cannot be challenged or turn out to be false. Expressives are, in this sense, performative. A representative example is (4.6), which I model on similar examples from Kaplan (1999)

(4.6) a. #That bastard Conner was promoted, if Conner is a bastard.
b. #If Conner is a bastard, then that bastard Conner was promoted.
c. That bastard Conner was promoted. #But probably he's not a bastard.

These results are fairly decisive. But they seem to tell us something rather different than the usual deniability tests do. The additional important fact is that expressives, unlike more familiar entailments, are also not challengeable by a hearer. Though discourses like (4.7) are perfectly well formed, they are not parallel to (4.8).

(4.7) a. Anne: That bastard Conner got promoted.
b. Kyle: Conner is not a bastard.

(4.8) a. Anne: Conner got promoted.
b. Kyle: Conner did not get promoted.

In (4.8), Kyle felicitously challenges Anne's assertion that Conner got promoted. Thus, the claim is the subject of debate; the proposition that Conner got promoted, offered by Anne, is not included in the common ground of these discourse participants. In (4.7), on the other hand, the contribution of the epithet *that bastard* as used by Anne does become part of the common ground. Kyle refuses to accept the characterization, but this does not in any way mitigate Anne's use of the epithet.

An essential step in coming to grips with this contrast is this: we analyze epithets and other expressives as expressing properties of speakers' emotional states.

In this way, we ensure that outright denials of their content by a hearer will make little sense. If we try to interpret Kyle's reply in (4.7) as a challenge to *Anne's* view of Conner, then the scenario becomes as far-fetched as one in which Kyle seriously challenges that Anne has a certain opinion or is in a state of pain. In all these cases, Kyle must deny an event that Anne has privileged access to.

4.3 Expressive adjectives and epithets

My primary sources of examples are *expressive attributive adjectives* (EAs) such as *damn* in (4.9a) and epithets like the expressions highlighted in (4.9b–d). Huddleston and Pullum (2002:36) and Soames (2002:57–58) identify the content of EAs as conventionally implicated; extending this insight to epithets systematizes existing observations about their interactions with commanding operators (Aoun et al. 2001).

(4.9)
a. Ed refuses to look after Sheila's <u>damn</u> dog.
b. Right after Chuck agreed to help out, <u>the jerk</u> boarded a plane for Tahiti.
c. Right after he agreed to help out, <u>that jerk Chuck</u> boarded a plane for Tahiti.
d. Every Democrat with [a proposal for reform]$_1$ claims [<u>the stupid thing</u>]$_1$ deserves public support.

We can characterize the content of these items as speaker-oriented (nonpresuppositional) assertions that are, in an intuitive sense, independent of the at-issue entailments. Each of these characteristics returns us to the original definition of CIs summarized in (4.3).

So EAs and epithets provide further data showing that definition (4.3) picks out a nonempty class of linguistic phenomena. A particularly important clause is (4.3c), which relativizes CI content to the speaker of the utterance. I offer two naturally occurring cases in which this is undoubtedly the intended interpretation of the EA:

(4.10) a. "We bought a new electric clothes dryer [. . .] Nowhere did it say that the <u>damn</u> thing didn't come with an electric plug!"[1]

b. "I remember practicing for my first Confession in the second grade and of course Sister role-played the priest. Trying to do a good job, I told her all the big sins [. . .] Never again!!!! For my Penance she made me say the <u>damn</u> rosary."[2]

These show that an embedded EA *can* be interpreted with widest-scope. They do not, however, show that it *must* be so interpreted. Three sorts of arguments fill this gap. The first appeals to standard presupposition holes (negation, modalization, conditionals, and questioning). Example (4.11) is representative:

(4.11) It's just not true that Sheila's damn dog is on the couch!

This sentence cannot be read as negating the speaker's disapprobation of Sheila's dog; it is it judged false if and only if Sheila's dog is not on the couch. Testing with the other holes reveals the same invariance, as seen in (4.6) above and also (4.12):

[1] <http://jjdavis.net/blog/arc20010325.html>

[2] <http://www.nunstories.com/SampleStories/SampleStories.html>

(4.12) a. I am not sure whether Conner is a jerk. #Is that jerk Conner coming to the party tonight?

b. It might be that Sue invited that jerk Conner. #But Conner might not be a jerk.

In all cases, expressive content scopes out of hole environments.

A second argument concerns tense, and marks a divergence between presuppositions and expressives. I observe first that presuppositions can (but need not) remain in the scope of tense operators, as seen in (4.13).

(4.13) Ed's dog died. (Hence, Ed does not at present have a dog.)

In contrast, expressives never end up in the scope of tense operators; the example in (4.14) is representative.

(4.14) That jerk Ed skipped work last week. #But Ed isn't a jerk now, not since he has started showing up regularly.

The expressive contribution of *jerk* is not interpreted in the scope of the past tense of the first sentence's main clause. The first sentence asserts that Ed is a jerk at the time of utterance, hence the incoherence of the continuation, which asserts that Ed is not a jerk at the time of utterance.

My third argument for the widest-scope (scopeless) nature of expressives derives from presupposition plugs. The generalization is the same as it is for tense, though the operators in question are different in that they are quite generally more likely to take n-ary tuples of propositions as their arguments, as discussed in chapter 6. In light of those arguments, it is especially important to check for embedding under plugs. Pairs like (4.15) fill out the picture.

(4.15) a. Clinton: The damn Republicans should be less partisan.
b. Bush: Clinton says the damn Republicans should be less partisan.

The sentence in (4.15b) is an unlikely report of Clinton's utterance (4.15a). Even those with a limited grasp of the language recognize that *damn*, even inside an indirect quotation, is heard as a contribution of the speaker of the utterance. Though *Clinton* is the subject of the propositional attitude verb in (4.15b), the content of *damn* is not relativized to Clinton, but rather to Bush, the speaker. The meaning of (4.15b) is roughly given by the pair of meanings in (4.16)

(4.16) a. Clinton says the Republicans should be less partisan
b. Bush looks with disapprobation upon Republicans

I refer to Quang 1971 and Cruse 1986:271–272 for the same generalization based on similar examples.

To report the content of *damn* in (4.15a), one must resort to a paraphrase (*Contemptuous of Republicans, Clinton says...*) or assign the EA a special intonation contour indicating that it is a quotative utterance: one signals this with heavy emphasis on the EA in speech, quotation marks in print; the result is subject to felicity conditions parallel to those of anaphoric resolution. It's worth stressing that, though quotative utterances can give rise to what appear to be embedded readings, a general analysis should treat them as scopeless. For instance, some speakers find, contrary to the judgment of Quang (1971), that *fucking* in (4.17) can be read as relativized to John's beliefs.

(4.17) John says that his landlord is a fucking scoutmaster.

Such non-speaker-oriented readings require heavy emphasis on *fucking*, an indication that they are quotative. In virtue of being in this quotative form, an EA can be attributed to an entity not mentioned in the sentence, as in (4.18).

(4.18) a. Sue: John's landlord is a fucking scoutmaster.
b. Eddie: Well, John wouldn't say that his landlord is a "fucking" scoutmaster. He rather admires scoutmasters, and so do I.

If these readings were a matter of scope, the EA in (4.18b) could be attributed only to John or the speaker, Eddie. Neither scoping gives us the intuitively correct reading. One might think that the addressee is another index we could exploit for (4.18). But suppose Eddie is speaking to a crowd. He reads Sue's claim aloud, then says (4.18b). Nothing changes about how we read "fucking". Thus, these observations stand as a third and important argument that EAs are syntactically embeddable but semantically unembeddable.

EAs are not alone among expressive modifiers in displaying this mix of properties. Epithets are also speaker-oriented even when embedded below propositional attitude verbs. I offer first an attested example; nowhere does the news story mentioned in (4.19) characterize the person who broke into Clements' house negatively.

(4.19) "The story says that the idiot broke into Clements' home and attacked and robbed him. Obviously, particularly since the guy only got $27, if Clements doesn't do something, the guy comes back and does it again a day or two later."[3]

Once again, we now know it is *possible* to interpret epithet content with widest-scope. Pairs like (4.20) permit us to strengthen this, assigning epithets to widest-scope status as well.

[3] <http://209.157.64.200/focus/news/780053/posts>. Both the news story and the reaction to it in (4.19) are at this page.

(4.20) a. Ellen: *Rambo* is a stupid movie.

b. Frank: I liked *Rambo*, but Ellen says the stupid thing isn't worth seeing.

This mini-discourse is strange. Though the context works to support a reading of the epithet in (4.20b) on which its propositional content is attributed to Ellen, we interpret it as an emotive contribution of Frank's (the speaker's). Aoun et al. (2001) make this observation about epithets in Lebanese Arabic, specifying that they should have a "main clause" interpretation (p. 386). As with EAs, apparently embedded readings reveal themselves to involve quotation. In (4.21b), the quotative "total snooze" is oriented neither to the speaker nor the matrix subject, but rather to Ellen, who is not directly mentioned in the sentence.

(4.21) a. Ellen: *The Godfather II* is a total snooze.

b. Frank: Well, Pauline Kael said that this "total snooze" is a defining moment in America cinema.

This quotative analysis seems right for the apparently embedded readings of epithets found by Kratzer (1999:(14)), which I provide in (4.22), and Schlenker (2003:(109)). Schlenker's cases, repeated in (4.23), are particularly useful: where the context supplies no agent to attribute the epithet to as a quotation, (4.23a), the speaker-oriented reading is the only one available, which here yields an emotive inconsistency.

(4.22) My father screamed that he would never allow me to marry that bastard Webster.

(Kratzer 1999:(14))

(4.23) a. $^{\#}$I am not prejudiced against Caucasians. But if I were, you would be the worst honky I know.

b. I am not prejudiced against Caucasians. But John, who is, thinks/claims that you are the worst honky he knows. (Schlenker 2003:(109a,b))

We can gather together the above observations at a slightly more technical level. If we ensure that expressive content is never the argument to an at-issue functor, then it will never end up in the scope of at-issue material. It will thus attain speaker-oriented status in the same manner in which main clause assertions attain this status. With an important qualification to allow, for example, **damn** to apply to **republican** in *damn Republicans*, we want at-issue and CI content to be impermeable to each other. Achieving this would have benefits not only in terms of scope, but also in terms of our ability to model the independence of the at-issue and CI dimensions (clause (4.3d)). At the level of expressive content, this would model speakers' intuitions that they can agree to any of the examples in (4.24) without committing themselves to the emotive baggage engendered by the modifiers.

(4.24) a. The damn Republicans are aggressively cutting taxes.

b. We saw that bastard Charlie at the pool hall.

$\mathcal{L}_{\mathrm{CI}}$ is just the sort of descriptive tool we need. At its very heart is the freedom to allow certain terms (those with a superscript c) to stay out of the scope of everything, thereby attaining main clause status.

The other clauses of Grice's (1975) definition (4.3) are also easy to match to facts about EAs and epithets. It is quite clear, for instance, that we are dealing with a specific group of lexical items; the content does not flow from the maxims and general considerations about how conversation works. (In the usual terminology, expressive content is ‘detachable’.) And, unlike presuppositions, EA content need not be entailed by the input context for felicitous use. A corollary of this is that we have additional support for the position that presupposition holes alone do not provide the tools we need to diagnose content as presupposed (Beaver 2001:19ff). One must also check the behavior of the content in question under tense operators, quantifiers (see below), and presupposition plugs.

4.3.1 An undistinguished syntax

A dominant theme of the above is the apparent widest scope of expressive content. This would easily be achieved in the absence of the CI hypothesis if it could be shown that the items in question had a syntax involving root-level adjunction, as discussed in chapter 5. The main purpose of this interlude in the semantic discussion is to head-off this alternative, by showing that neither EAs nor epithets display syntactic properties that suggest a nonstandard syntax. A more concrete result of this attention to the syntax is that it suggests an ideal shape for the interpreted parsetrees, and in turn highlights inevitable deviations from that ideal (see (4.33)).

English has a rich array of EAs (*bleeding*, (*gol*)*darn*, *mother loving*, and so forth), and new ones are coined fairly regularly, often by popular media as substitutes for swear words. It is useful to use *damn*, *darn*, and *frigging* in examples, because they are unambiguously EAs (the regular adjective is *damned*, and *frigging* is a tame alternative to *fucking* but without the literal meaning for most speakers).

Huddleston and Pullum (2002:553) show that EAs are syntactically much like other strictly attributive adjectives (e.g., *former*, *premier*). They are restricted to prenominal position, but freely intermingle with other adjectives:

(4.25) a. Sheila said that we must look after her (biggest) friggin' brown dog.

b. "What's the Big Friggin' Deal About Sony PlayStation 2?"[4]

In German, where attributive adjectives are marked for case, EAs are not distinguished from other adjectives in this sense:

(4.26) a. "Du hast kein verdammtes Wort gesagt."
*you have no.*ACC *damn.*ACC *word said*
'You didn't say a damn word.'[5]

b. "dass die verdammte Industrie zu geizig ist, [...]"
*that the.*NOM *damn.*NOM *industry too miserly is* [...]
'that the damn industry is too miserly [...]'[6]

The case-markers on the adjectives and determiners are the same as those found with regular attributive adjective constructions.

There are apparently no restrictions on the kind of determiner that can head a nominal containing an EA:

(4.27) a. The company says that every damn piece of software we use has to be made by them!

b. Ed claims that no damn idea of his should be ignored!

This is an important point. As discussed in the next section, EAs, though nominal-internal, often modify the entire proposition expressed by the immediate clause. Similar properties are found with adjectives like *occasional* in *An occasional native strolled by*

[4] <http://www.zdnet.com/anchordesk/stories/0,10738,2645040,00.html>

[5] Leonard, Elmore. 2000. *Schnappt Chili*. Translated by Hans M. Herzog. München: Wilhelm Goldmann Verlag (p. 146).

[6] <http://members.chello.at/manfred.mayer2/elektronik.htm>

(Stump 1981), which means the same thing as *Occasionally, a native strolled by.* One might seek to extend to EAs Zimmermann's (2000) syntactic movement analysis, on which the infrequency adjective raises to form a quantifier $INFREQ$, denoting a family of sets of event–individual pairs. At least three arguments suggest that this is not a fruitful direction in which to head. First, as Zimmermann shows (p. 295), infrequency adjectives permit adverbial readings only with articles and possessives. The adverbial reading of EAs is not limited by the determiner. Second, infrequency adjectives are required to appear adjacent to the determiner for their adverbial readings; in contrast, both examples in (4.25) can involve the EAs as clause-level modifiers, but neither is determiner-adjacent. Finally, EAs express no notion of (in)frequency; $INFREQ$ is quite obviously not the proper denotation for these expressions.

The upshot of the above syntactic discussion is that an EA plays no special role in the syntax of a nominal it appears in, beyond simply adjoining as any modifier would. That is, EAs determine routine structures of the form in (4.28), in which they are simply left-adjoined modifiers (node labels highly negotiable).

(4.28)

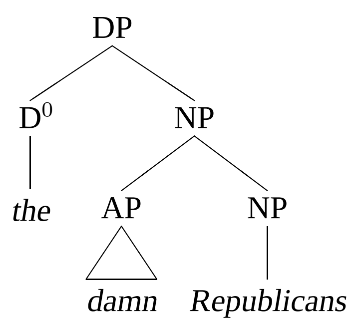

It seems safe to conclude that the contrasts between EAs and other attributive adjectives don't follow from properties of the structures they determine.

We can say the same for epithets. The work of Jackendoff (1972:§4.1), Lasnik

(1976), Aoun and Choueiri (2000), and Aoun et al. (2001) shows that epithets are cross-linguistically much like full nominal expressions in their syntactic distribution. For instance, they are sensitive to c-commanding antecedents, as in (4.29).

(4.29) $^{\#}\{\text{Paul}_1/\text{No musician}_1/\text{He}_1\}$ thinks [the vain snob]$_1$ is tiresome.

Where discourse considerations remove this c-command effect, as in (4.30a), epithets are allowed. I provide (4.30b) to show that we can have variable binding between the relevant positions, which strongly suggests that they are in a c-command relationship.

(4.30)
a. "In 1654 a friend had written him$_1$ to ask if Pascal$_1$ could solve the *problème des parties*, or problem of points."[7]
b. The professor wrote every student$_1$ to ask if the lazy bum$_1$ could solve the problem of points.

One feature of epithets is noteworthy: Aoun and Choueiri (2000:2–3) report that Lebanese Arabic epithet nominals containing names are unique in taking an extra definite marker: *ha-l-habiile Sami* transliterates as 'this the idiot Sami', where the initial demonstrative is a kind of dummy. Though I do not pursue this line here, one might regard the extra definite determiner as a morphological marker of the semantic differences between epithets and regular definite nominals. One could extend this idea to English, in the form of a more abstract extra definiteness layer.

4.3.2 Lexical meanings

Our goal is to separate the expressives' content from that of the regular at-issue assertion, so it is no surprise that the rule of **CI application**, repeated in (4.31), plays a

[7] Bass, Thomas A. 1985. *The Eudaemonic Pie*. Boston: Houghton Mifflin Company (p. 117).

central role.

(4.31) **CI application**

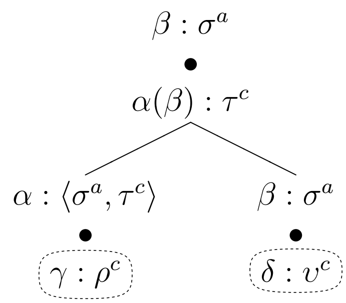

Recall that optional material is inside dotted lines, and that the structures do not contain a linear adjacency relation at all, so the left-to-right order in the diagram is technically arbitrary. The bullet, •, is a metalogical device for keeping independent lambda terms clearly separated.

It is this rule that permits us to have subtrees such as (4.32), in which the meaning of *damn* applies to its sister to form a CI proposition, and that sister's term is passed up to the mother as well.

(4.32)

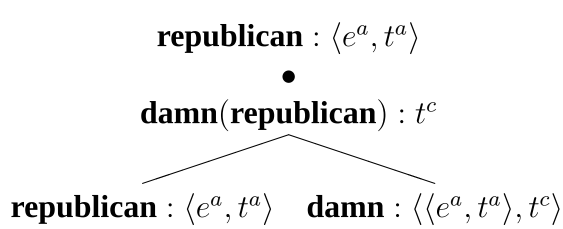

The question now is just what function or functions *damn* names. This is of course a special case of the broader question of how to write lexical entries for EAs. They are a rich, open class. But semantically, the members seem not to be distinguished from each other. Though there is a kind of scale, with *darn* as the tame end and *fucking* at the obscene end, the details are heavily discourse-conditioned: a playful use of *fucking* might sound less angry or disapproving than a stern *damn*. I simplify by giving only the narrow semantics of EAs, mapping them all to the same lambda term.

But which term? Another issue that I have not addressed thus far is the fact that EAs need not be interpreted as taking their common noun sisters as arguments. The immediately containing full noun phrase or the entire clause can also be targets:

(4.33) a. We have to look after Sheila's damn dog.

b. Nowhere did the instructions say that the damn machine didn't come with an electric plug!

With (4.33a), the speaker probably does not express disapprobation of all dogs, but rather just Sheila's; (4.33b) arguably expresses the speaker's frustration with the fact that the machine in question arrived plugless. Since the syntactic evidence militates against movement of attributive adjectives, but the existence of these readings indicates that some can act as clause-level functors, it seems safe to conclude that this does not happen via syntactic processes. The treatment of sentences as pairs of trees, one of them a semantic parsetree, lets us model this essential semantic fact (it is a question about function–argument structure) without messing with the syntax at all. We simply allow that in these cases, the syntactic and semantic parsetrees have different shapes. The structure in (4.34) is an example.

(4.34)

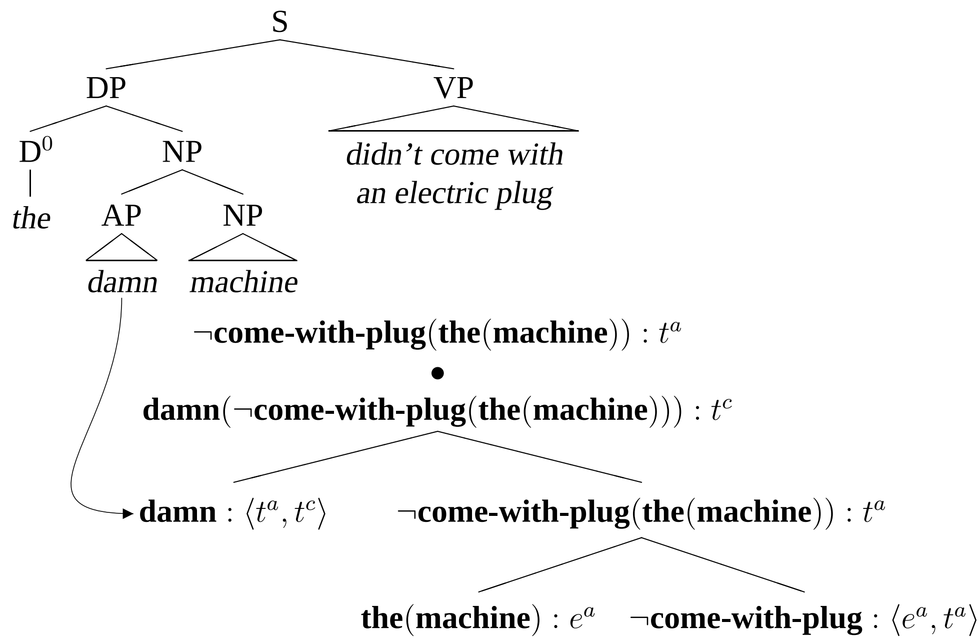

At the level of denotations, the variability of the arguments to an EA indicates polymorphism in the domain of the EA meaning. I offer a general lexical entry, on which an EA can take any argument in $\langle \tau^a, t^a \rangle$ to produce a term of type t^c:

(4.35) $$\left\{ \begin{array}{r} \textit{damn} \\ \textit{bloody} \\ \vdots \\ \textit{fucking} \end{array} \right\} \rightsquigarrow \lambda X.\, \mathbf{bad}({}^{\cap}X) : \langle \langle \tau^a, t^a \rangle, t^c \rangle$$

The nominalizing type-shifter ${}^{\cap}$ is that of Chierchia (1984). When defined extensionally, it takes any function and returns the plural individual composed of all members of the input set. (In symbols, ${}^{\cap} = \lambda X.\, \iota x[\forall y[X(y) \leftrightarrow y \leqslant x]] : \langle \langle \sigma^a, t^a \rangle, e^a \rangle$, where $\leqslant$ is the 'part of' relation and ι is the definite operator. See also example (3.27), page 136.)

The translation in (4.35) contains a function called **bad**. It seems likely that there is also a positive correspondent to this. John Kingston (p.c., 2/03) notes that that British

English adjective *brilliant* has an EA semantics that expresses the speaker's positive attitude. The meaning for elements in this class is presumably as in (4.35), but with **good** in place of **bad**. However, any investigation of the exact meaning of these modifiers is bound to be seriously complicated by the fact that any sufficiently strong word can lose its meaning, or see that meaning reversed. It might, therefore, be better to regard **bad** as the name of the function that says, roughly, 'the speaker is in a heightened emotional state regarding X'. In any event, to be more precise about the specifics of these meanings would imply a degree of understanding of the semantics of EAs that I do not possess. My interest is in managing the content, whatever it is. The important thing is the typing of the EA denotation: it takes $\langle e^a, t^a\rangle$ inputs and returns propositional CI results. This suffices to account for EAs' inability to appear in predicative position; I illustrate using a *be* that takes properties into same, but the result holds across theories of the copula.

(4.36)

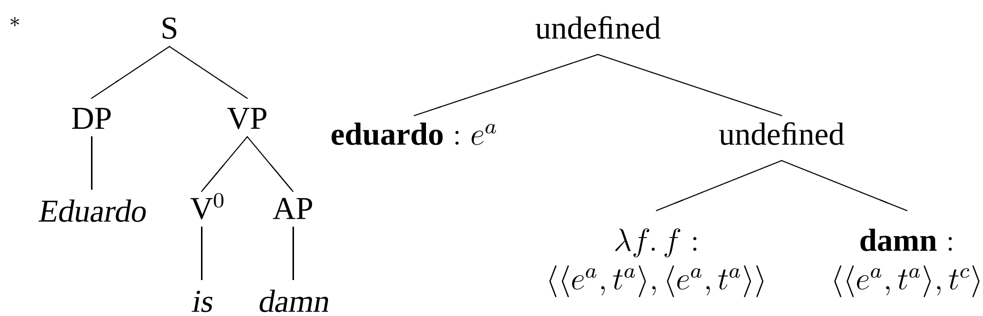

In (4.36), *be* denotes an identity function on properties. Since the problem is one of typing, the result holds equally if *be* translates as $\lambda f \lambda x.\, f(x) : \langle\langle e^a, t^a\rangle, \langle e^a, t^a\rangle\rangle$. Other theories of *be* also fail to produce a meaningful meaning; though we might find cases in which *be* is a function on generalized quantifier types ($\langle\langle e^a, t^a\rangle, t^a\rangle$), these could not be functions on generalized quantifier *conventional implicature* types, for the simple reason that we have no types whose first members are CI types — the definition of the set of types, repeated in (4.37), which organizes and delimits the possible translations for natural language expressions, does not contain types with CI inputs.

(4.37) a. e^a, t^a, and s^a are basic at-issue types for $\mathcal{L}_{\mathrm{CI}}$.
b. e^c, t^c, and s^c are basic CI types for $\mathcal{L}_{\mathrm{CI}}$.
c. If σ and τ are at-issue types for $\mathcal{L}_{\mathrm{CI}}$, then $\langle \sigma, \tau \rangle$ is an at-issue type for $\mathcal{L}_{\mathrm{CI}}$.
d. If σ is an at-issue type for $\mathcal{L}_{\mathrm{CI}}$ and τ is a CI type for $\mathcal{L}_{\mathrm{CI}}$, then $\langle \sigma, \tau \rangle$ is a CI type for $\mathcal{L}_{\mathrm{CI}}$.
e. If σ and τ are at-issue types for $\mathcal{L}_{\mathrm{CI}}$, then $\langle \sigma \times \tau \rangle$ is a product type for $\mathcal{L}_{\mathrm{CI}}$, a subset of the set of at-issue types for $\mathcal{L}_{\mathrm{CI}}$.
f. The full set of types for $\mathcal{L}_{\mathrm{CI}}$ is the union of the at-issue and CI types for $\mathcal{L}_{\mathrm{CI}}$.

Functional CI types appear only by means of clause (4.37d), which does not allow a CI to appear as the input type. CIs are strictly output types.

This same logic affords an account of why EAs, despite forming a loose scale of strength, are not gradable. Since gradable modifiers like *very* are functions from at-issue meanings into same, they cannot apply to the CI dimension of a word like *damn*. The parsetrees for examples like (4.38) are not licensed; see (4.38e)

(4.38) a. *That's a {quite/very/really/super} damn dog.
b. *You have no idea just how {damn/fucking} Sally's dog is.
c. *Pico wrote a more damn novel than Brio wrote a play.
(cf. Kennedy and Merchant 2000:(7a))
d. *Juan is too fucking Swedish to know about jet skiis.
e. undefined

very : $\langle\langle d^a, \langle e^a, t^a\rangle\rangle, \langle d^a, \langle e^a, t^a\rangle\rangle\rangle$ **damn** : $\langle\langle e^a, t^a\rangle, t^c\rangle$

Here, I quietly extend $\mathcal{L}_{\text{CI}}$ with a type for degrees; we can assume that the domain of d^a is the set $\mathbb{R}$ of real numbers on their natural ordering. The details are not significant; as with the copular clauses above, the heart of the explanation is that we don't have terms that take input CI terms. Hence, there is no way that $\alpha(\textbf{damn})$ can be well-formed, for any choice of terms α.

We do allow an alternative structure. The adverbial can adjoin not to the adjective but to the modified common noun. This generally corresponds to something grammatical. It seems that in this case we do not have the EA applying to the gradable adjective, but rather functioning in propositional (attitude-indicating) capacity. I therefore invoke the **isolated CIs** rule in the description:

(4.39) a. Sue's dog is really fucking mean.

b.

$\textbf{really}(\textbf{mean}) : \langle d^a, \langle e^a, t^a \rangle\rangle$

- $\textbf{really} : \langle\langle d^a, \langle e^a, t^a \rangle\rangle, \langle d^a, \langle e^a, t^a \rangle\rangle\rangle$
- $\textbf{mean} : \langle d^a, \langle e^a, t^a \rangle\rangle$
  - $\textbf{fucking} : t^c$
  - $\textbf{mean} : \langle d^a, \langle e^a, t^a \rangle\rangle$

I repeated the rule of **isolated CIs** in (4.40), both because it is essential to (4.39) and because it plays a significant role in the analysis of EAs, which seem often to avoid all interaction with at-issue material.

We saw motivation for (4.40) in chapter 3, section 3.6.1, in the form of appositives like *The students, most* $($*of them*$)$ *sophomores, were unhappy with the assignment.*

(4.40) **isolated CIs**

$$\beta : \tau^a$$

$$\alpha : t^c \qquad \beta : \tau^a$$

$$\bullet$$

$$\gamma : \rho^c$$

This rule allows for readings of EAs on which they seem not to take any arguments at all, as in the (non)interaction with *really* above. Such EAs function to express the speaker's disposition; the parsetree in (4.39) is interpreted as a pair of propositions: the at-issue proposition that Sue's dog is really mean, and the CI proposition that the speaker is in a heightened emotional state (probably anger).

I note that the structure in (4.38e) is also in fact a licit one if the EA appears in its propositional form. This permits the meaning of *very* to move up unmodified, as it does in (4.39). It then takes the non-EA adjective as its argument. The result is well formed only if this argument is gradable.

The lexical entries — both the functional one in (4.35) and the propositional one assumed for the gradable modifier cases — facilitate an explanation for why CIs are unembeddable. But the general result is largely independent of individual meanings. In order to embed a CI meaning under an at-issue operator A, it would have to be the case that A had an extensional type of the form $\langle t^c, \tau^a \rangle$. But we have no such types. As a result, there is no provision in the set of local tree conditions for taking CI content into at-issue content. An example helps bring out this property of the logic: if we take the parsetree in (4.34) and embed it under *Ed says*, we have the following semantics:

(4.41)

$$
\begin{array}{c}
\textbf{say}(\neg\textbf{come-with-plug}(\textbf{the}(\textbf{machine})))(\textbf{motorola}) : t^a \\
\quad \textbf{motorola} : e^a \qquad \textbf{say}(\neg\textbf{come-with-plug}(\textbf{the}(\textbf{machine}))) : \langle e^a, t^a \rangle \\
\qquad \textbf{say} : \langle t^a, \langle e^a, t^a \rangle \rangle \qquad \neg\textbf{come-with-plug}(\textbf{the}(\textbf{machine})) : t^a \\
\bullet \\
\textbf{damn}(\neg\textbf{come-with-plug}(\textbf{the}(\textbf{machine}))) : t^c
\end{array}
$$

By our provision for interpreting trees, this is defined as having the value $\langle 1, 1 \rangle$ iff Ed says the machine doesn't come with a plug, and it is bad that the machine doesn't come with a plug. The speaker-orientation of the latter, CI proposition follows from the same interpretive specifications that make the former proposition speaker-oriented.

It seems quite likely that EAs are just a special group of modifiers with only expressive realizations. When one looks at adjectives that express subjective judgments, one finds that they alternate between at-issue and CI based readings. I noted one such example in chapter 1, in (1.20). I repeat that example here:

(4.42) Edna is at her friend Chuck's house. Chuck tells her that he thinks all his red vases are ugly. He approves of only the blue ones. He tells Edna that she can take one of his red vases. Edna thinks the red vases are lovely, selects one, and returns home to tell her housemate,

"Chuck said I could have one of his lovely vases!"

We find here an obligatory mismatch between the sister to *lovely* in the syntax and the sister to **lovely** in the semantics:

(4.43)

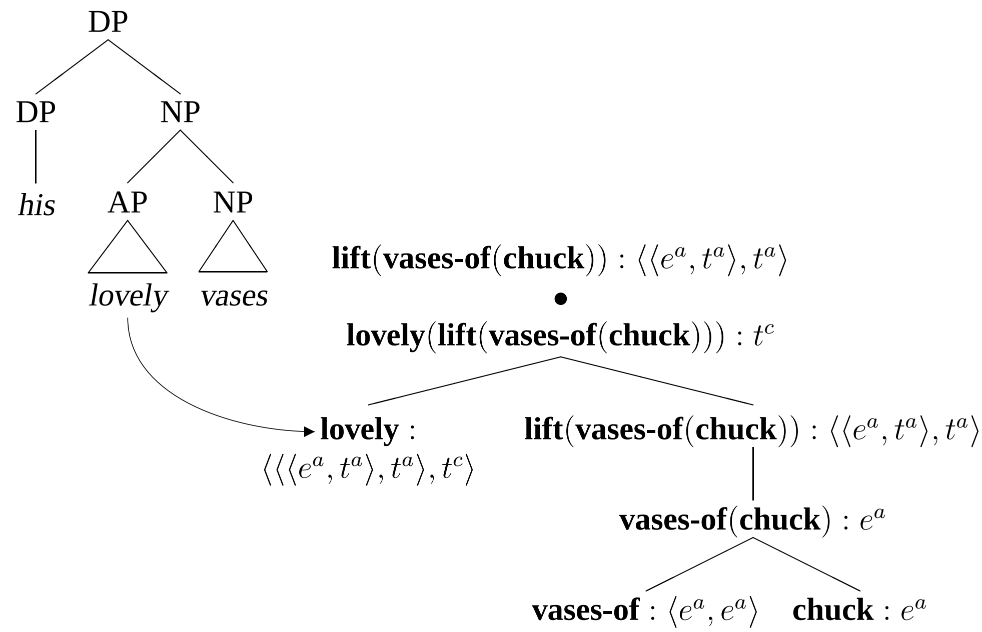

Here, **lift** is a feature term, taking entities to their generalized-quantifier correlates. We could also have its converse, **lower**, taking generalized quantifiers to entities (Partee 1987).

This analysis represents a reading on which the speaker characterizes only Chuck's vases as lovely. She might find other vases to be hideous looking. The proper meaning is obtained if we assume that, in its CI guise, **lovely** has the same type as an EA: it takes functional things as arguments but shifts them to their entity-level denotations in order to say something about them. This parallel typing is not necessary, but it does unify the two different kinds of adjective rather neatly.

It is more challenging to find lexical denotations for epithets, since they have the same rich internal structure as regular noun phrases. I propose that all epithets have the structure of those that place an appositive modifier on a name (Huddleston and Pullum 2002:447–448). Where the name is absent, a free variable fills its spot:

(4.44) a. {that/the} stupid jerk Eddie

b. {that/the} stupid jerk x_{25}

A fuller picture is given in (4.45).

(4.45)

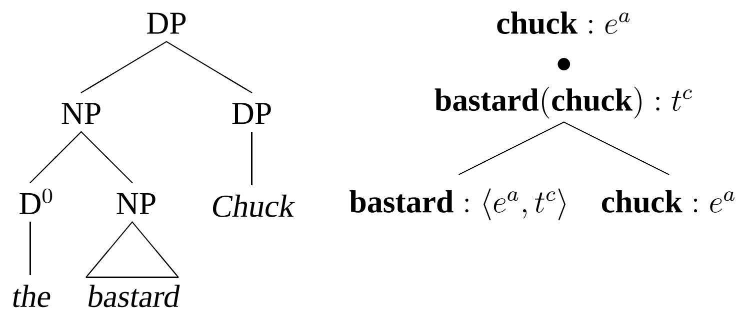

I note that this analysis ignores the definite article in the semantics. This seems in line with the fact that, unlike regular definites, epithets do not presuppose that a unique entity meets the conditions specified by their descriptive content.

As I present it here, the analysis depends upon a free variable. But, broadly speaking, we find no motivation for (free) variables in this domain. The above could be recast in terms of bound variables and in turn in terms of combinators. Where we have saturated expressions denoting open propositions, a variable-free analysis would have property-denoting expressions in both the CI and at-issue dimensions. The details of how to do this are spelled out by Jacobson (1999:134–135).

4.4 Quantifiers and a variable environment dimension

So far I have considered only epithets that have a referential semantics. However, we have already seen cases in which an epithet appears to be (in some sense) dependent upon a higher quantifier. In (4.46a), I repeat example (4.30b) from above; the example is modelled on the attested case (4.30a).

(4.46) a. The professor wrote every student$_1$ to ask if the lazy bum$_1$ could solve the problem of points.

b. The judge told every dead-beat dad$_1$ that the bum$_1$ must help out.

Aoun et al. (2001) also report instances of quantifier–epithet connection:

(4.47) kəll muttahame saʔalto ʔəza ha-l-maʒduube nħabasit
*each suspect.*SF *asked.*2P *whether this-the-idiot.*SF *imprisoned.*3SF
'Each suspect, you asked whether this idiot was imprisoned.'
(Aoun et al. 2001:373, (6))

(4.48) kəll muttahame tfeeʒaʔto lamma / laʔanno ʕrəfto ʔənno
*each suspect.*SF *surprised.*2P *when because know.*2P *that*
ha-l-maʒduube nħabasit
3*-the-idiot.*SF *imprisoned.*3SF
'Each suspect, you were surprised when/because you knew that this idiot was imprisoned.'
(Aoun et al. 2001:388, (18a))

A large part of Aoun et al. 2001 is devoted to explaining why such relationships are apparently not possible in Lebanese Arabic unless an island boundary intervenes. Their explanation for why (4.49a) is impossible is that it would have to have a logical form like (4.49b).

(4.49) a. * kəll muttahame ʕrəfto ʔənno ha-l-ma-ʒduube nħabasit
*each suspect.*SF *know.*2P *that this-the-idiot.*SF *imprisoned.*3SF
'Each suspect, you know that this idiot was imprisoned.'
(Aoun et al. 2001:(5a))

b. x is an idiot; each suspect x is such that you know x was imprisoned

But, given only what they say, it is mysterious why (4.47) is not bad because its logical form contains an unbound variable:

(4.50) x is an idiot; each suspect x is such that you were surprised $\{$when/because$\}$ you knew that x was imprisoned

It would seem that the present account is in the same bind. Based on the referential cases, a natural interpretation of the components of meaning for (4.46b) would seem to be (4.51).

(4.51) a. at-issue: the judge told every dead-beat dad x that x must help out
b. CI: x is a bum

If this is the correct representation, then we need to adjust the logic somehow. In virtue of the fact that (4.51a) is a distinct term from (4.51b), the variable x in the CI meaning is free; it gets its interpretation from the assignment.

But attention to the semantics of the epithet's expressive content in these cases reveals that it would be a mistake to bind the variable in the CI. The connection between the two dimensions of meaning is not nearly so tight as this would imply. The analysis we seek is actually one on which the propositional content of the epithet is a generic quantification over the restriction on the at-issue determiner — a looser connection between the two dimensions than genuine variable binding. Let's see why this is so.

A factual argument against binding across the dimensions comes from the insensitivity of the epithet's content to changes in the at-issue quantifier it connects with. The expressive content of (4.52) is identical to that of (4.46b) despite changes in the at-issue meaning resulting from variation in the quantified object of *tell*.

(4.52) a. The judge told almost $\{$every/no$\}$ dead-beat dad that the bum must help out.
b. The judge told $\{$few/many/most$\}$ dead-beat dads that the bums must help out.

For all examples in (4.52), the CI is identical to the CI in (4.46b): a generic quantification of roughly the form 'generally, dead-beat dads are bums'. We would wrongly predict variation in the nature of the CI if the CI contained a variable bound from the at-issue dimension. (This is another point of contrast with presuppositions, which display a rather complex mix of properties when in the nuclear scope of a quantifier (Heim 1983; Cooper 1983:152–154; Krahmer 1998:§4).)

These facts are admittedly subtle. It would be helpful to supplement them with factual considerations of a less delicate nature. One argument of this form concerns the well-known generalization, discussed in chapter 3, section 3.6.1, that only referring expressions can associate with nonrestrictive modifiers (Karttunen 1976; McCawley 1998:451, Potts 2002a:83; and others), a generalization that extends even to instances in which the anchor is a bound variable (hence locally referential). I repeat example (3.90) here:

(4.53)

a. Every student_1 spoke with a psychiatrist of hers_1 that welcomes housecalls.

b. *Every student_1 spoke with a psychiatrist of hers_1, a caring individual who welcomes housecalls.

c. Sally_1 spoke with a psychiatrist of hers_1, a caring individual who welcomes housecalls.

Epithets are nonrestrictive; (4.46b) is not equivalent to (4.54).

(4.54) The judge told every dead-beat dad who is a bum that he must help out. $\neq$ (4.46b)

On the present analysis, the nonrestrictiveness of epithets follows directly from the fact that their descriptive content cannot possibly influence the at-issue proposition expressed. The descriptive fact alone provides independent reason to doubt that the CI aspect of epithets is directly in the scope of a quantifier.

But, as noted, there is an undeniable link between the at-issue quantifier and the epithet's content. A factually accurate analysis of (4.46b) is (4.55). (G is a generic quantifier.)

(4.55) The judge told [every dead-beat dad]$_1$ that [the bum]$_1$ must help out.

a. at-issue: $\forall x[\textbf{db-dad}(x) \rightarrow \textbf{tell}(\textbf{must}(\textbf{help}(x)))(x)(\textbf{the-judge})]$

b. CI: $\mathsf{G}x[\textbf{db-dad}(x) \rightarrow \textbf{bum}(x)]$

The technical question is how to get **db-dad**, the restriction to dead-beat dads, into the CI dimension in a systematic way. The restriction logic of Gawron (1996) provides exactly the needed mechanism. The aim is to work towards a translations like (4.56).

(4.56) a. environment: $x \mid \textbf{db-dad}(x)$

b. meaning: $\forall x[\textbf{tell}(\textbf{must}(\textbf{help}(x)))(x)(\textbf{the-judge})] : t^a$
$\mathsf{G}x[\textbf{bum}(x)] : t^c$

The environment serves to restrict all possible values of the variable x to those that have the property named by **db-dad**. An innovation of restriction logic is making this restriction hold even in the scope of later quantifiers; the semantics for quantifiers like $\forall$ and G appeal to the information in the environment for their restrictions. The result is an identity for x beyond its binding quantifier. The environment is itself a new dimension of meaning.

I close this section by noting a stubborn bit of context sensitivity: referential epithets involve predication of their descriptive content of some individual, whereas dependent epithets place this descriptive content in the nuclear scope of a generic quantification over an independently provided restriction. Thus, whereas a referential *the bum* x_{25} has its CI dimension given by $\textbf{bum}(x_{25}) : t^c$, a dependent use has the CI dimension $\mathsf{G}x[\textbf{bum}(x)] : t^c$. It would be good to remove this context-dependency, but I do not see how. I note, though, that it is not the only contrast of this type: (4.53) indicates that regular pronouns behave differently depending on how they receive their values.

4.5 A scope-shifting alternative

It is worth heading off an alternative analysis that attempts to locate the unusual properties of expressive content modifiers in the theory of scope-shifting, thereby reducing expressive content to regular at-issue meanings. A reduction of this sort is assumed to be feasible in Kaplan 1989:55, fn. 71, a short footnote in which Kaplan characterizes epithet examples like *John says that the lying S.O.B. who took my car is honest* in terms of what he calls 'pseudo de re'. He suggests that quantifying-in is a suitable mechanism for handling such examples. That is, he treats what I have called the CI content of epithets in terms of standard scope-shifting mechanisms.

It is a bit foolhardy to develop an approach along these lines in any detail, since I do not advocate it and no one else really has either to date. It seems worthwhile nonetheless, because it is important to make clear that a description of expressives is of course not outside the bounds of intensional logic. Intensional logic, particularly the sort that employs explicit, object-language abstraction over intensional indices, is an extremely powerful tool that admits of lots of extensions. (Quite often, the indices are *quintuples and growing*: world, time, place, speaker, addressee,. . .) But one can still challenge the approach. I mount a two-pronged attack: (i) the scope-shifting solution must be augmented by a variety of different metalogical conditions to prevent overgeneration; and (ii) it must call upon the central notions of the $\mathcal{L}_{\mathrm{CI}}$ description — lexical marking and multidimensionality — to achieve descriptive adequacy. All told, the $\mathcal{L}_{\mathrm{CI}}$ description appears to be the cleaner of the two.

To amplify these points, I first show how we would describe examples like Kaplan's in a language without CI types and conditions like **CI application**. Here is a variation on the above-mentioned example from Kaplan 1989:

(4.57) a. John says that the liar is honest.

b. $[\![\mathbf{say}_{@}(\lambda w.\, \mathbf{honest}_{w}(\mathbf{the}(\mathbf{liar}_{@}))(\mathbf{john})]\!]^{\mathcal{M}}$

= 1 iff every belief world w for John in @ is such that the unique liar in @ is honest in w

I use @ as a free world variable representing the actual world. (Because considerations of time and space are not immediately relevant, I assume that this index denotes a member of the set of possible worlds.)

Of course, we would have to stipulate that expressives are always evaluated at the actual world index @. That is, we would have to saturate their intensional arguments *in the lexicon*. Perhaps this is a reasonable stipulation to make, though I have my doubts that we should make the lexicon dependent upon the structures in this way. The important thing for present purposes is that this would amount to lexical marking of expressives. In $\mathcal{L}_{\mathrm{CI}}$, we achieve lexical marking simply by appeal to the type theory. This decision has many advantages, chief among them that we can rely on quite general facts about lambda calculi to derive the restrictions we need. There is no need for ad hoc stipulations to cut down the well-formed expressions of the logic to those in which expressives have a particular world variable as their intensional argument.

This is by no means the only stipulation necessary. Much more must be said before this approach can account for the various restrictions we meet easily in $\mathcal{L}_{\mathrm{CI}}$. Copular clauses provide a representative example. Recall that the type theory for $\mathcal{L}_{\mathrm{CI}}$ and a motivated denotation for *damn* is able to account for the restriction exemplified in (4.36), repeated in (4.58a).

(4.58) a. *Eduardo is damn.

b. $\mathbf{be}_{@}(\mathbf{damn}_{@})(\mathbf{eduardo})$

The translation in (4.58b) meets the condition that the expressive is evaluated at the actual world index. So an additional statement would be needed to block just this kind of example.

The nonrestrictiveness of expressives is also a problem for this approach. As noted above, epithets and EAs are never restrictive. But if we interpreted their content (in the actual world) as we interpret the content of regular adjectives, we would allow them to function restrictively. I illustrate the conundrum in (4.59), which is best considered alongside (4.60).

(4.59) a. The damn Republicans succeeded. (#Luckily, relatively few Republicans deserve my disapprobation.)

b. $\mathbf{succeed}_{@}(\mathbf{the}_{@}(\mathbf{damn}_{@}(\mathbf{republican}_{@})))$

(4.60) a. The red apples are inedible. (Luckily, relatively few of the apples are red.)

b. $\mathbf{be}_{@}(\mathbf{inedible}_{@})(\mathbf{red}_{@}(\mathbf{apple}_{@}))$

The second of these examples seems to be described correctly using the intensional language. The continuation indicates that *red* might have a restrictive semantics: the red apples might be a proper subset of the full set of apples. But EAs cannot be restrictive, as indicated by the infelicitous continuation in (4.59a). However, the translation could well involve applying the meaning of *damn* to the meaning of *Republicans* to return a set that is properly contained in the set of Republicans. Similar considerations hold for epithets, as well as for the expressives discussed elsewhere in this chapter. To eliminate the possibility of restrictive interpretations, we would require meaning postulates or their equivalent — again, metalogical conditions that actually restrict the logic to one that is smaller than that given by the unadorned definition of the well-formed formulae. I venture no objections to meaning postulates per se here (see example (3.151), page 202). I simply stress that they are unnecessary for these facts if we are working in $\mathcal{L}_{\mathrm{CI}}$.

My final objection to this approach is the most fundamental of them all. The issue is multidimensionality. This dissertation offers numerous arguments that expressives (and supplements) have a multidimensional semantics. But the more general intensional-logic translations in this section are one-dimensional: they would be interpreted as individual propositions (however those are construed). Thus, even if we were to fully flesh out the conditions described above, we would still be faced with the task of fundamental revisions to the logic. A sentence like *The damn Republicans won* cannot denote a single meaning. It must denote at least a pair of them. The bottom line is that in order to model speakers' intuitions that expressive content is a comment upon the at-issue core, the at-issue account needs to shift the domain of sentences from $\{0, 1\}$ to $\{0, 1\}^n$, the set of all ordered tuples with elements drawn from $\{0, 1\}$. This multidimensionality is the defining feature of the CI account. Accepting it is in large part a concession that we need a logic like $\mathcal{L}_{\mathrm{CI}}$, rather than the usual sort of one-dimensional intensional language. I fully expect that revisions to $\mathcal{L}_{\mathrm{CI}}$ will prove necessary in the future. I am highly skeptical that these revisions will move us towards something like Montague's IL. CIs indicate that a departure from that logical space is necessary.

Since I began this section with a description from Kaplan 1989, it seems proper to close by pointing out that Kaplan (1999) drastically revised his assessment of epithets as posing no special difficulties. The newer paper begins from the premise that, in Kaplan's (1999) words, "It may be that the *primary* problem in semantics is not what does this or that *mean*, but rather in what *form* should we attempt to say what this or that means" (p. 3). He then uses this premise to begin building a semantic theory that is guided by expressives. That is, where Kaplan (1989) sought to assimilate expressives to regular semantics, Kaplan (1999) seeks to assimilate regular semantics to expressives.

4.6 Honorifics in Japanese

Kaplan (1999) makes an explicit connection between honorification and expressives. In addition, he homes in on just what makes honorifics special:

(4.61) "In sum, in addition to the desire to be *held* in respect, people desire to be *paid* respect and honorifics can be the coin of that payment." (Kaplan 1999:28)

This nicely articulates the performative aspect of honorific marking: the act of using it is in itself an act of paying respect. Honorifics are, in this sense, the positive counterparts to things like epithets and EAs, which tend to be (though are not invariably) negative.

The class of honorifics divides in two. The better-studied class contains those that involve honorification of the meaning of a syntactic argument of a predicate. I henceforth refer to these as *argument-oriented* honorifics. Harada (1976) uses 'propositional honorifics' for this class. I avoid that terminology because I think all honorifics contribute propositions of one kind or another.

Argument-oriented honorifics are realized as bound morphemes on nominalized verbs. They indicate something about the speaker's relationship to one of the verb's arguments. For instance, example (4.1a) above, repeated in (4.62), brings along the honorific contribution that the speaker regards Yamada as socially superior to him. This meaning is independent of the at-issue content in a way that we have seen many times already.

(4.62) **Japanese verbal (subject) honorification**

Sensei-wa	eigo	ga	o-wakari-ni	nar-u
the-teacher-SUBJ	*English*	NOM	HON-*understanding*-DAT	*become*-IMP

'The teacher understands English.' (Toribio 1990:535, (1a))

The second group contains the *performative honorifics*. These are less well studied, and their nature is somewhat controversial. Toribio (1990:535, fn. 2) writes that "it is unclear whether performative honorification exists at all". Since there are multitudes of examples in Harada 1976 and elsewhere, this seems intended as an expression of skepticism about the idea that performative honorifics are properly grouped with argument-oriented honorifics. This is a question I leave open. What I can show is that performative honorifics admit of the same sort of semantic treatment as utterance-modifying adverbs like *frankly*, with which they share a host of properties.

What is not in doubt is that the two classes of honorific have considerably different distributions. I therefore address them in separate subsections, beginning with the argument-oriented variety.

4.6.1 Argument-oriented honorification

We have available to us a rich literature on the syntax of honorifics in Japanese (e.g., Harada 1976; Toribio 1990; Boeckx and Niinuma 2003). There is also sensitive sociolinguistic work on the issue of when honorifics are felicitously deployed. But theoretical linguists have been shy about addressing the meaning of honorifics. As noted above, Harada (1976) is skeptical that such a description falls within the purview of theoretical linguistics. I repeat (4.2) in (4.63a), along with a second, more specific injunction.

(4.63) a. "The question as to what properties are associated with a person who is referred to through honorifics is not a problem to which a grammatical description is addressed, though it is an interesting matter for sociolinguistic researches."

(Harada 1976:500)

b. "The notion of social superiority, on the other hand, does not seem definable in a simple, culture-independent way."

(Harada 1976:501)

This section presents a dissenting opinion regarding (4.63a). It seems likely that honorifics will push us in new directions regarding the models for linguistic theory. But will they really push model-theoretic techniques past the breaking point? This seems unlikely. The justification that (4.63b) offers is not persuasive: the meanings of words like *love* and *respect* are culture-bound, yet semanticists are happy to treat these as two-place relations. If it turns out that the conditions for inclusion in these relations are based on astrological signs, chemicals in the brain, or political-party affiliations, this will not place such notions outside the bounds of formal-modelling techniques, or even call for adjustments to existing methods.

The same seems true for honorifics. The informal descriptions we have all converge on a single notion for the contribution of the subject honorifics in (4.64)–(4.65).

(4.64) Sasaki sensei wa watasi ni koo o-hanasi-ni nat-ta.
Sasaki sensei TOPIC *I* I.O. *this way speak* BE
'Sasaki *sensei* told me this way.' (Harada 1976:(2a))

(4.65) Yamada sensei-ga o-warai-ni nat-ta.
Yamada teacher-NOM HON-*laugh*-DAT *be*
'Professor Yamada laughed.' (Shibatani 1978:54, cited in Toribio 1990:539)

The honorific is realized morphologically on the main verb, which is nominalized (Toribio 1990); a transliteration of (4.65) might read 'Yamada *sensei* (honorable) laugher became'. Its contribution is the proposition that the denotation of the subject, *Sasaki sensei*, is socially superior to the speaker. The description is easily transferred to a model-theoretic setting as a particular two-place relation. Call this relation **honorific**, the 'socially superior to' relation. A semantic analysis of (4.65) should converge on the pair of meanings in (4.66).

(4.66) a. at-issue: $\mathbf{laugh}(\mathbf{yamada}) : t^a$
b. CI: $\mathbf{honorific}(\mathbf{the\text{-}speaker})(\mathbf{yamada}) : t^c$

The appositive-like expression *sensei-ga* might contribute a second CI, the proposition that Yamada is a professor (or, simply a respectable person; the exact contribution is context-dependent; Harada 1976:509).

In honorifics we find a variation on a theme of EAs: their syntactic argument appears not to be their semantic argument. In (4.65), the honorific appears on the verb, yet the target of honorification is the meaning of the subject. In addition to such subject-oriented honorifics, there are object- and oblique honorifics. All appear on the verb. How do we get the meanings where we need them to be?

My answer is heavily dependent upon the syntactic literature, which seems to have converged on the hypothesis that honorification involves a kind of feature-checking. Thus, in an example like (4.65), we arguably have a syntactic structure in which the subject and verb are linked. The existing proposals of this form call upon rather elaborate assumptions about the look of the structures (e.g., Toribio 1990:539, (12)). I choose to abstract away from many of the details, offering a structure like (4.67).

(4.67)

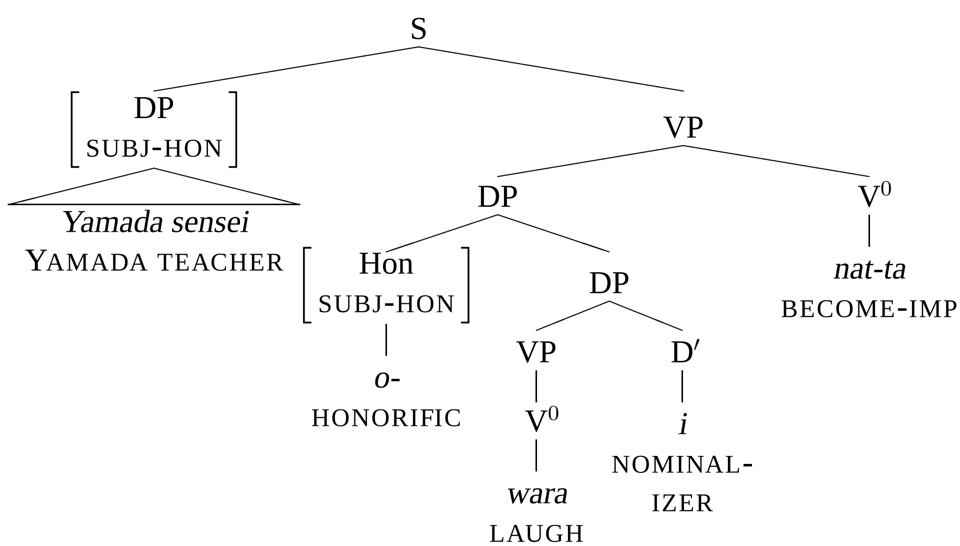

The required statement is that the HON feature on the subject is the semantically potent one; for morphological or syntactic reasons, the morphology appears on the verb rather than the subject, but this is incidental to the parsetrees. The semantics for (4.67) is thus (4.68).

(4.68)

$\textbf{laugh}(\textbf{yamada}) : t^a$

$\textbf{Yamada} : e^a$ $\quad$ $\textbf{laugh} : \langle e^a, t^a \rangle$

$\bullet$

$\textbf{honorific}(\textbf{Yamada}) : t^c$

$\textbf{honorific} : \langle e^a, t^c \rangle$ $\quad$ $\textbf{Yamada} : e^a$

Honorifics of this form are embeddable, but retain their widest-scope semantics even when placed inside indirect quotation environments. Thus, example (4.69) does not attribute to John the knowledge that the speaker regards Mary's mother as socially superior. (My thanks to Shigeto Kawahara , (p.c., 1/03), for the example and insights into its meaning.)

(4.69) John-wa [Mary-no okaasan-ga koojyoo-de hataraite-rassharu
John-TOPIC *Mary*-GEN *mother* *factory*-.LOC *work*-HON
koto-wo] shitteiru.
fact-ACC *knows*
'John knows that Mary's mother, who is socially superior to me, works in a factory.'

Controlled infinitival complements offer an opportunity to see what happens when there is multiple honorific marking with the same content. Harada (1976:546–547) reports that, for examples equivalent to *Honorable Yamada decided to go to Karuizawa*, in which the argument to the meaning of the embedded predicate is Yamada, there is a preference for honorific marking on the matrix verb alone. But marking on just the complement is also possible, though it is "a bit awkward and less polite" (p. 547). Marking both the matrix and the complement verbs is "somewhat too polite to use in ordinary honorific contexts" (p. 547). An example of double-marking is given in (4.70), with honorific verb forms in boldface.

(4.70)

Yamada	sensei	wa [	Karuizawa	ni	**o-ide**	**ni**	**nar-u**	]$_S$	koto	ni
Yamada	*teacher*		*Karuizawa*		HON-*go*	DAT	*become*-IMP		COMP	
o-kime	**ni**	**nat-tu**.								
HON-*decide*	DAT	*become*								

‘Professor Yamada decided to go to Karuizawa.’

(*matrix and embedded honorifics; Harada 1976:546*)

The sense of redundancy that inheres in such examples is probably formally akin to the redundancy one senses when a speaker chooses to indicate his emotional state with repeated uses of EAs or other expressive markers of that negative sort. Consider, for example, the rather excessive string of EAs in (4.71a), alongside the closest nonexpressive counterpart to it, (4.71b).

(4.71) a. The *friggin'* dog is on the couch. Take the *friggin'* animal and put it outside. If you can't keep the *friggin'* thing off the couch, we're sending it to Siberia.

b. # *I dislike the dog.* It is on the couch. *I dislike the dog.* Take it outside. *I dislike the dog.* If you can't keep it off the couch, we're sending it to Siberia.

The first is tiresome. The second is not really a coherent discourse due to its redundancy. This contrast no doubt traces back to the utterance-oriented nature of expressives. Since the utterance situation is in near constant flux, there is much less a chance of redundancy than there is with regular at-issue assertions about described situations.

4.6.2 Performative honorifics

As noted, performative honorifics (‘polite speech’), exemplified in (4.72), are different from argument-oriented honorifics in a number of ways that impact how we treat them in terms of the description logic.

(4.72) Ame ga huri-masa-ta.
rain SUBJ *fall*-HON-PAST
'It rained.' (Harada 1976:502)

Harada (1976) writes that "one uses performative honorifics in order to talk 'politely' to the addressee, to make one's speech sound 'milder' " (p. 507). Like utterance-modifiers, performative honorifics do not permit any kind of syntactic embedding; Harada says that "the few complement constructions that do permit performative honorifics to occur are interpretable, without exception, as 'direct discourses' " (p. 544). An illustrative example:

(4.73) *Boku wa [kyoo Yamada sensei ga **ki -mas-u** / **o-ide ni nari-**
I SUBJ *today Yamada teacher* TOP HON *come* / *come*
mas-u] koto o sukkari wasure-te i-ta.
COMP OBJ *entirely forget*-GERUND *be*
'I completely forgot that Professor Yamada is coming today.'
(Harada 1976:544)

The basic characterization of performative honorification, along with its restriction to root-clauses, suggest that the upper layer of our discourse structures holds the key to an understanding of their contribution. It seems that a performative honorific modifies the utterance relation: the speaker utters the sentence in question politely. The description is based on the meaning in (4.74), which is a term drawn from $\mathcal{L}_U$, the discourse language.

(4.74) $\ulcorner\textbf{perf-hon}\urcorner \stackrel{def}{=} \lambda S. \ulcorner\textbf{politely}\urcorner(\ulcorner\textbf{utter}\urcorner(S)(\ulcorner\textbf{the-speaker}\urcorner) : \langle u, t\rangle$

This is a function from full sentences to truth values. I assume that the performative morphology invokes this extra layer of meaning, as in (4.75).

(4.75) $\ulcorner$**perf-hon**$\urcorner$($\ulcorner$Ame ga huri-masa-ta$\urcorner$) : t

$\ulcorner$**perf-hon**$\urcorner$: $\langle u, t\rangle$ $\quad$ $\ulcorner$Ame ga huri-masa-ta$\urcorner$: u

$\Big\langle$ S [*Ame ga huri-masa-ta*] , **fall**(**rain**) : t^a [**fall** : $\langle e^a, t^a\rangle$ $\quad$ **rain** : e^a] $\Big\rangle$

The dashed line represents the interpretation function. It takes us from the constant of $\mathcal{L}_U$, $\ulcorner$Ame ga huri-masa-ta$\urcorner$, to the pair of trees that provides this constant with its interpretation. The term $\ulcorner$**perf-hon**$\urcorner$($\ulcorner$Ame ga huri-masa-ta$\urcorner$) : t is interpreted as 1 if the speaker said "Ame ga huri-masa-ta" politely, else it is interpreted as 0. We determine the value of the sentence the speaker uttered by taking advantage of the interpretation rules for $\mathcal{L}_{\text{CI}}$. Thus, a performative honorific contributes the same sort of meaning as an utterance-modifying adverb, viz., that the speaker is saying something in a particular way.

4.7 German Konjunktiv I

German has a fairly well-studied system of voice-marking, Konjunktiv I, employed primarily to indicate that the speaker wishes to distance himself from the propositional content expressed. Konjunktiv I is, for this reason, the dominant mode for reporters to use in news stories. The following example is typical (I gloss Konjunktiv I forms with 'KONJ'):

(4.76) Fritz glaubt, dass Maria krank sei.
*Fritz believes that Maria sick is.*KONJ
'Fritz believes that Maria is sick.' (von Stechow 2002:(125c))

A speaker of (4.76) disavows himself of any commitment, even via implicatures, to the proposition that Maria is sick. Tellingly, when we have a first-person matrix subject with a present-tense matrix verb, the Konjunktiv I becomes deviant. Compare:

(4.77) a. #Ich glaube, dass Maria krank sei.
*I believe that Maria sick be.*KONJ
'I believe that Maria is sick.' (von Stechow 2002:(125a))

b. Ich glaubte, dass Maria krank sei.
*I believed that Maria sick be.*KONJ
'I believed that Maria is sick.' (von Stechow 2002:(125b))

When the matrix verb is in the past, as in (4.77b), the example is felicitous: the speaker conveys with such examples that his beliefs have changed.

Factives predicates do not permit the subjunctive in their complements:

(4.78) #Fritz ärgerte sich darüber, dass Maria krank sei.
*Fritz annoyed self there-over that Maria sick be.*KONJ
'Fritz was annoyed about the fact that Maria was sick.'
(von Stechow 2002:(126))

The descriptive generalization that we should capture, indicated by the above facts and speakers' intuitions about basic cases, is (4.79).

(4.79) Use of the Konjunktiv I in a clause C with content p indicates that the speaker is not publicly committed to the truth of p.

One should note well the limited claim the Konjunktiv I makes. Whereas EAs are either strongly negative (*fucking*) or strongly positive (*brilliant*), Konjunktiv I is cool, detached. It does not indicate that the speaker is committed to the negation of the propositional content in question, only that he has not made any moves that commit him to the truth of this content. Of course, use of a factive entails such a commitment. With (4.79) formulated, we can see why (4.78) goes wrong:

(4.80) #Fritz ärgerte sich darüber, dass Maria krank sei.
*Fritz annoyed self there-over that Maria sick be.*KONJ
'Fritz was annoyed about the fact that Maria was sick.'
(von Stechow 2002:126)

a. (4.80) is defined only if Maria is sick
b. at-issue entailment of (4.80): Fritz is annoyed that Maria is sick
c. CI of (4.80): the speaker is not committed to the proposition that Maria is sick

The presupposition (4.80a) and the CI contradict each other. These facts are easily described in $\mathcal{L}_{\text{CI}}$. The first step is to posit a lexical entry for the Konjunktiv I marking. I henceforth name the relevant functor **kI**. It is defined as in (4.81).

(4.81) $$\mathbf{kI} \overset{def}{=} \lambda \mathbf{p} \lambda w.\, \exists w'[w \mathsf{E} w' \wedge \neg \mathbf{p}(w')] : \langle\langle s^a, t^a\rangle, \langle s^a, t^c\rangle\rangle$$

I use E to name the relation of epistemic possibility. Thus, the formula $w\mathsf{E}w'$ holds just in case the pair consisting of the world named by w as its first member and the world named by w' as its second member is in the set **E**. We can gloss $w\mathsf{E}w'$ with 'w' is an epistemic alternative for w'. The denotation of **kI** simply indicates that there are epistemically accessible worlds in which the input proposition is false. Since we evaluate all formulae of $\mathcal{L}_{\text{CI}}$ inside the model associated with the speaker (see chapter 2, section 2.8), **kI** contributes the equivalent of (4.79).

The meaning of **kI** requires it to appear at the clause level in parsetrees, another deviation from the syntactic form. A typical structure is represented in (4.82).

(4.82) $\textbf{believe}(\textbf{sick}(\textbf{maria}))(\textbf{fritz}) : \langle s^a, t^a \rangle$

- $\textbf{fritz} : e^a$
- $\textbf{believe}(\textbf{sick}(\textbf{maria})) : \langle e^a, \langle s^a, t^a \rangle \rangle$
  - $\textbf{believe} : \langle \langle s^a, t^a \rangle, \langle e^a, \langle s^a, t^a \rangle \rangle \rangle$
  - $\textbf{sick}(\textbf{maria}) : \langle s^a, t^a \rangle$ • $\textbf{kI}(\textbf{sick}(\textbf{maria})) : \langle s^a, t^c \rangle$
    - $\textbf{kI} : \langle \langle s^a, t^a \rangle, \langle s^a, t^c \rangle \rangle$
    - $\textbf{sick}(\textbf{maria}) : \langle s^a, t^a \rangle$

The interpretive rule **parsetree interpretation** determines that (4.82) denotes a pair of sets of worlds: the set of all worlds in which Fritz believes that Maria is sick, and the set of worlds in which the speaker is open to the falsity of the proposition that Maria is sick.

This is, as far as I know, a new treatment of Konjunktiv I. The rest of this section is devoted to comparisons between this account and existing ones. The first, due to Schlenker (2003) and von Stechow (2002), seeks to classify Konjunktiv I marking as a *logophoric tense*, building on what is known about logophoric pronouns, which are found overtly in Ewe, and seem also to be trigged by emphatic reflexives and *own* in English. A defining feature of a logophoric item on these accounts is that it must appear inside an attitude report; see Clements 1975:171 and Schlenker 2003:59ff for full definitions.

In the hands of Schlenker in his section 5.1.2, and von Stechow (2002), the classification corresponds to a statement in the grammar of the following form (see von Stechow 2002:43 and Schlenker 2003:76 for their particular statements):

(4.83) A verb marked with Konjunktiv I must have its world argument supplied by a different world than the one that saturates the matrix verb's world argument.

Using the notation of section 4.5 above, we can say that the condition in (4.83) blocks (4.84a), while allowing (4.84b).

(4.84) a. $\textbf{believe}_{@}(\lambda w.\, \textbf{sick}_{@}(\textbf{maria}))(\textbf{the-speaker})$

b. $\textbf{believe}_{@}(\lambda w.\, \textbf{sick}_{w}(\textbf{maria}))(\textbf{the-speaker})$

The $\mathcal{L}_{\mathrm{CI}}$ description does not require an extra statement of this form. The contribution of **kI** has the effect of ensuring that the complement to the higher attitude verb is not evaluated with widest-scope (i.e., at the actual world index). The result of such a widest-scope evaluation would be the contradiction in (4.85).

(4.85) $[\![\textbf{believe}_{@}(\lambda w.\, \textbf{sick}_{@}(\textbf{maria}))(\textbf{the-speaker}) : t^{a}]\!]^{\mathcal{D},s,a}$

$=$ 1 iff every epistemic alternative w in @ for s is such that Maria is sick in @

$[\![\exists w[@\mathsf{E}w \wedge \neg\textbf{sick}_{w}(\textbf{maria})] : t^{a}]\!]^{\mathcal{D},s,a}$

$=$ 1 iff there are epistemic alternatives to @ in which Maria is not sick

Since all content is relativized to the speaker, the first meaning asserts that the speaker believes Maria is sick in the actual world, whereas the second says that the speaker believes there are epistemically accessible worlds in which Maria is not sick. This is a pathological mental state. Thus, basic conditions on consistent assertions ensure that if Konjunktiv I marking appears on a clause, then that clause is not evaluated at the actual world index. This would put the CI and the at-issue entailment in conflict, as in (4.85).

Schlenker's (2003) proposal differs markedly from von Stechow's (2002). Schlenker offers a felicity condition. Like von Stechow, Schlenker is concerned with morphosemantic details that I think do not play a role in the present discussion. So I offer the following greatly simplified version of his condition (Schlenker 2003:86, (86)):

(4.86) a. Let CG, a set of worlds, represent the context of utterance. Indicative marking on a clause C is defined only if the world of evaluation for C is in CG.

b. Konjunktiv I is used only if indicative marking "*in the same Logical Form* would result in a presupposition failure".

Thus, Schlenker sets up a competition model. Konjunktiv I appears only where the indicative cannot. As he says, "in a nutshell, whenever it is presupposed that s(w) is in the Common Ground, the use of $-$indicative^{2*} [roughly: indicative marking —C. P.], and hence of Konjunktiv I, should be precluded" (p. 86).

There are inherent drawbacks to adopting a competition model of this sort, as it radically alters the logic itself. Consider the semantic parsetrees of $\mathcal{L}_{\mathrm{CI}}$, comparable to the Logical Forms that Schlenker adopts as a meaning language. The parsetrees are singly-rooted, connected structures representing (the meanings for) individual sentences. All the tree-admissibility conditions limit the shape of these objects individually. There is no way to enforce a condition like (4.86b), because it requires access to sets of parsetrees. Therefore, to make sense of it, we need to redefine the theory so that its objects are sets of parsetrees. Determining the size and nature of these sets would be an important and difficult matter. If competition models can be avoided, then they should be avoided.

For extensive elaboration of this point, I refer to Potts and Pullum 2003. The context of that paper is phonological theory, but all the same theoretical considerations apply in syntax and semantics.

However, even in the context of a competition model — even if we take the radical step of adjusting grammatical theory in order to accommodate such principles — a serious factual problem confronts (4.86b): it is often the case that Konjunktiv I marking is optional. Both of the following are grammatical, though the first might be preferred in many cases:

(4.87) a. Fritz glaubt, dass Maria krank sei.
*Fritz believes that Maria sick be.*KONJ

b. Fritz glaubt, dass Maria krank ist.
*Fritz believes that Maria sick be.*INDICATIVE

Both examples mean that Fritz believes that Maria is sick. The second could be preceded or followed by an assertion that Maria is in fact sick, or by an assertion that she is not sick. The first could appear felicitously only with the speaker's denial that Maria is sick. That is, while the second might, but need not, convey the conversational implicature that the speaker believes Maria is sick, the second is prevented, by the Konjunktiv I marking, from conveying such a conversational implicature. The very existence of both alongside one another, with no narrow semantic difference between them, indicates that the competition principle is, at the very least, not stated accurately.

The presuppositional alternative is also an important contender. To get us started, I first offer, in (4.88), a basic presuppositional denotation for Konjunktiv I. The subscript p distinguishes this term as the $p$resuppositional or $p$artial analysis of Konjunktiv I.

(4.88) $\mathbf{kI}_p \rightsquigarrow \lambda \mathbf{p} \lambda w \ \ddagger \exists w'[w \mathsf{E} w' \wedge \neg \mathbf{p}(w')] \ddagger \ . \mathbf{p} : \langle\langle s^a, t^a\rangle, \langle s^a, t^a\rangle\rangle$

The expression between the doubled daggers effects a precondition for definedness. Many similarities between (4.81) and (4.88) are evident. First, (4.88) is an identity function on at-issue propositions. Though this is not built into (4.81) directly, it is handled by the CI logic: since (4.81) is a CI term, any argument to it is passed on unmodified. Second, (4.81) translates the action of (4.81) into a definedness condition — a precondition for truth. These substantive parallels severely limit our ability to tease apart the analyses. Where defined, the presuppositional meaning in (4.88) matches the action of (4.81) exactly.

I think we can nonetheless locate stumbling blocks for (4.88). Much of von Stechow's (2002) discussion is concerned with the contrast in (4.89).

(4.89) a. Ich dachte, ihre Yacht sei länger als sie ist.
*I thought your yacht be.*KONJ *longer than it is*
'I thought your yacht was longer than it is.'

b. #Ich dachte, ihre Yacht sei länger als sie sei.
*I thought your yacht be.*KONJ *longer than it be.*KONJ
'I thought your yacht was longer than it was.'

For von Stechow, Konjunktiv I marking in the comparative is ungrammatical because it forces that comparative to be interpreted inside the scope of *thought*. The result asserts that all of the speaker's belief worlds w are such that the length of the addressee's yacht in w is greater than its length in w, which is nonsense.

The $\mathcal{L}_{\mathrm{CI}}$ description provides much the same explanation. Konjunktiv I marking on the verb in this example means that the speaker is not committed to an actual length for the yacht. This means that we must interpret the clause internal to the belief context, else the result would be a contradiction: the speaker would be committed to a certain (range of) lengths by the at-issue semantics, but he would deny such a commitment with the Konjunktiv I marking. Thus, the comparative is interpreted inside the belief context, where the explanation takes the same form as it does for von Stechow.

But we differ in our predictions for multiply-embedded cases such as (4.90a), which is von Stechow's example (124), with his judgment.

(4.90) a. ? Fritz behauptet, Maria glaube, meine Yacht sei länger als
*Fritz maintains Maria believe.*KONJ *my yacht be.*KONJ *longer than*
sie tatsächlich sei.
*it in-fact be.*KONJ
'Fritz maintains that Maria believes that my yacht was longer than it in fact was.'

b. All of Fritz's belief worlds w_f in @ are such that all of Maria's belief worlds w_m in w_f are such that the speaker's yacht is longer in w_m than it is in w_f

The question von Stechow addresses with this case is whether it can have the indexing in (4.90b). He writes, "I don't get this reading, but better examples might exist" (p. 43).

I am also unsure of the facts. The difficulty could stem from the inclusion of the adverb *tatsächlich* ('in fact'), which is plausibly analyzed as requiring its complement to be evaluated in the actual world. That is, even when embedded, this factuality operator might demand widest-scope interpretation. This would conflict with the requirements of the Konjunktiv I marking. In the context of the $\mathcal{L}_{\mathrm{CI}}$ description, the result would be a contradiction: **in-fact** would impose primary scope commitment to truth, whereas **kI** would impose primary scope lack of commitment to truth.

The important question for present purposes is whether or not the semantics of Konjunktiv I can ever be embedded. A potentially relevant example is something like (4.91).

(4.91) a. Ich weiss, dass Maria krank ist.
I know that Maria sick is
'I know that Mara is sick.'

b. Fritz behauptet, dass Alonzo glaubte, dass Maria krank sei.
*Fritz maintains that Alonzo believes that Maria sick be.*KONJ
'Fritz maintains that Alonzo believes that Maria is sick.'

The question is whether (4.91b) could be felicitously preceded (4.91a). The $\mathcal{L}_{\text{CI}}$ description predicts markedness for any discourse containing these two sentences. No matter how deeply embedded the Konjunktiv I marking, it invariably contributes the expressive proposition that the content of the marked clause is not something the speaker is committed to. The feature-based accounts of von Stechow (2002) and Schlenker (2003) leave more room for scope variability. I am not at present in a position to resolve the factual question that could decide among these accounts. So I note only that if Konjunktiv I is embeddable, then a multidimensional account would remain viable. We could employ the product types of $\mathcal{L}_{\text{CI}}$ for this purpose (see chapter 1, section 1.4.5, and chapter 6).

4.8 Conclusion

I mentioned in section 4.5 that Kaplan (1989) uses the term 'pseudo de re' for, roughly speaking, epithets. His description suggests that he would welcome EAs under this heading as well. So it is worth providing the following snippet from his footnote:

(4.92) "I do not see that the existence of the *pseudo de re* form of report poses any issues of theoretical interest to make it worth pursuing."
(Kaplan 1989:555–556, fn. 71)

I hope to have made a convincing case that this dismissal is too hasty. EAs and epithets help validate Grice's (1975) definition of CIs as speaker-oriented comments upon the at-issue core of utterances. The account also suggests a method for managing expressive content.

Chapter 5

The supplement relation: A syntactic analysis

5.1 Remarks

Early in chapter 3, I offered various arguments in favor of an integrated semantics for supplements. Facts concerning linear adjacency, case-marking, nonextraposability, and correlations between supplements and other kinds of adjunction all quite strongly suggest that supplements are modifier constructions, adjoined in a familiar way and forming constituents with their anchors. Once this syntax is adopted, the conventional implicature (CI) logic is essential to deriving the proper meanings.

However, given the nature of past theoretical work on supplements, it seems wise, before fixing the CI premise, to look closely at a major competing hypothesis. I call this the *supplement relation* analysis, because it is usually based on the idea that supplements adjoin to structures via a designated supplement relation that is disjoint from dominance. McCawley (1982, 1987, 1989, 1998) was a long-time advocate of analyses with this general form; it forms the basis for the descriptions of Huddleston and Pullum (2002); and its essential components are recoverable from the work of Emonds (1976) and Culicover (1992). The analysis captures what is right about the reoccurring *wide-scope conjunction* analysis of (Lakoff 1966) without stumbling on the well-known arguments against that position (Ross 1967:§6.2.4; Pullum 1979:§4.1.1; Green 2000; and, from a much different perspective, Loetscher 1973). One might think that, working from this premise, we could trace the special behavior of supplements back to the structures, with the semantics following as a result of a compositional interpretation. The CI hypothesis would be dispensable.

My overarching conclusion is that the supplement relation leaves us far from a satisfactory treatment of the syntax and semantics of supplements. The supplement relation is not motivated by the constructions under the discussion here — it arguably obscures important generalizations — and the intuition guiding the wide-scope conjunction analysis is better cashed out in terms of a CI logic. However, the class of supplements and supplemental-type expressions is enormous and surely quite heterogeneous. It may well be that the supplement relation is essential to modelling some of them properly. Hence, the work I do in this chapter of fleshing out and honing the supplement-relation analysis might prove useful in other studies of supplemental expressions. But not the ones that are the focus of this chapter and the previous one.

This discussion proves useful for a higher-level reason as well. As I said in chapter 1, the study of CIs has suffered in part because few have made attempts to formalize them, to provide them with substance. This dissertation is a preliminary attempt, but I hope others follow it. The supplement relation is a partial alternative to $\mathcal{L}_{\mathrm{CI}}$. But what we find is that the supplement relation "alternative" in fact duplicates, in natural language syntactic terms, the essential ingredients of the CI analysis: lexical marking and multidimensionality. Since the novel tree geometry it imposes does not have motivated application in the area of expressive modifiers (chapter 4), it is likely that we would need the CI logic anyway even if we added this new dimension to our syntax. It might be that the supplement relation is useful elsewhere. But it seems that it will work in tandem with a CI logic, rather than serving as a replacement for it.

There are many variants of the supplement relation analysis, some of which incorporate features of the wide-scope coordination analysis that I believe dates to Lakoff's (1966) work on supplementary relatives. Rather than address each analysis individually, which would be rather tedious, I instead base the discussion on the analysis described by McCawley (1998:§13b). This analysis actually contains in it, in some form, the central features of all the analyses I am familiar with. Thus, this basis provides ample opportunity to address alternatives along the way.

5.2 McCawley's (1998) analysis

Lakoff (1966) proposes that supplementary relatives involve transformational derivations that have them adjoined to the root node in the interpreted structure. Thus, (5.1a) is interpretively equivalent to (5.1b).

(5.1) a. The officer arrested Clyde, who was the subject of a long manhunt, before he could strike again.

b. The officer arrested Clyde before he could strike again and Clyde was the subject of a long manhunt.

Effectively the same analysis is described by McCawley (1998:§13b). Once we abstract away from the details of the individual transformational grammars Lakoff and McCawley adopt, we see that the only difference is in the question of whether there is an overt *and* in the interpreted structure. For McCawley, the coordinator is *who*; the derivations are given as follows by McCawley (1998:449,(20)):

(5.2) a.

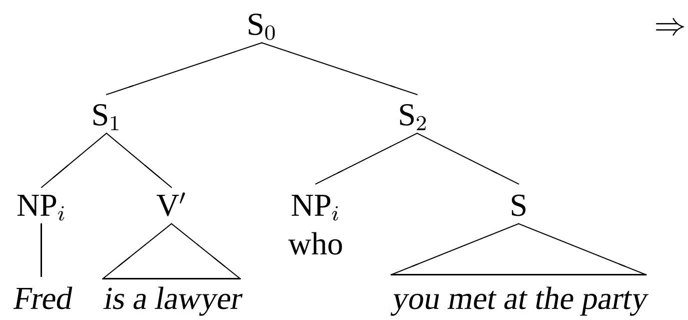

b.

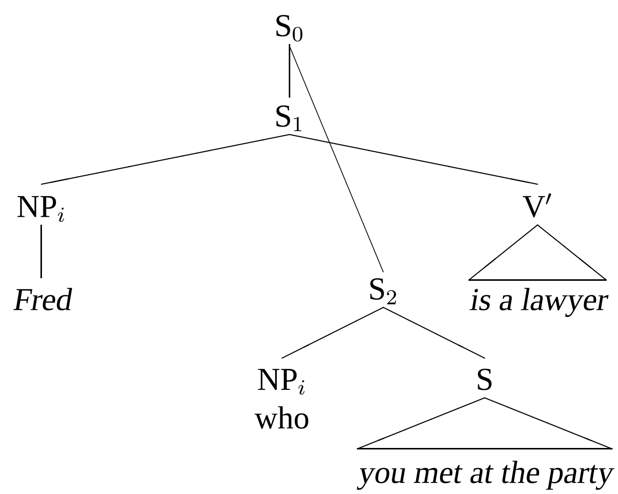

The analysis has three main parts: the coordinate interpreted structure, the mapping to the surface structure, and the nature of the surface structure itself. The next three subsections review each part in turn, addressing both factual and technical objections at each stage. My primary descriptive aim is to show that coordination and supplementation are not substantively related, a task I undertake by questioning the assumption that (5.2a) is anything like the structure that we interpret. My primary theoretical aim is to show that even if we adopted this kind of analysis, we would still need to appeal to something quite like $\mathcal{L}_{\mathrm{CI}}$ to achieve a grammar of supplementation.

5.3 The coordinate interpreted structure

I address first those issues that surround the interpreted structure, since the problems in this area are well known. For Lakoff, McCawley, and most others who have looked at these facts, the interpreted structure is an abstract underlying one. But nothing I say hinges on moving in this direction; the abstract coordination could be a logical form derived from the surface form, as in Aoun et al.'s (2001) analysis of Lebanese Arabic epithets as abstract main clause coordinations. The claim I dispute is the close connection between coordination and supplementation that this kind of analysis depends upon.

It did not take long for researchers to find fault with the coordination analysis of supplements. Ross (1967:§4.2.4.3) noted that coordinating two clauses of differing illocutionary force generally results in ungrammaticality. He went on to show, in his §6.2.4.1, that no such markedness arises when the main clause and supplement containing it have differing illocutionary force (see also Pullum 1979:§4.1.1). The examples in (5.3) indicate the negative correlation between coordination and supplementation in this area.

(5.3) a. Did the officer arrest Clyde, who was the subject of a long manhunt, before he could strike again?

b. *Did the officer arrest Clyde before he could strike again, and {Clyde/he} was the subject of a long manhunt.

c. *Clyde was the subject of a long manhunt, and did the officer arrest {Clyde/him} before he could strike again?

d. Clyde was the subject of a long manhunt. Did the officer arrest {Clyde/him} before he could strike again?

(5.4) a. Do the dishes, which are waiting for you in the sink!

b. *Do the dishes, and {the dishes/they} are waiting for you in the sink.

c. *The dishes are waiting for you in the sink, and do {the dishes/them}.

These examples suggest a general argument. We see that it is generally not possible to coordinate clauses of differing illocutionary force. Supplementary relatives (and most others) offer deemphasized assertions. But assertive supplements appear in clauses of varied illocutionary force, with no sign of awkwardness. Since the issues here seem to trace back to interpretation, it is a mistake to assimilate supplements to coordination in the interpreted structure. True, one might object that the restrictions in (5.3) are not narrowly grammatical, but rather arise from principles of cooperative conversation. In support, one might cite cases like (5.5).

(5.5) a. It'll be very hot, so take plenty to drink.
(Huddleston and Pullum 2002:1332, (26ia))

b. They've finished the job, but why did they take so long?
(Huddleston and Pullum 2002:1332, (26ib))

c. It certainly looks very good, but isn't it rather expensive?
(Huddleston and Pullum 2002:1732, (4iii))

The objection is cogent; it is unclear how a hearer should react to (5.3): with agreement, or with a salient answer? In the exceptions in (5.5), the declarative that comes first offers content that is either already known or not expected to be controversial. Hence, a hearer knows how to react to the second conjunct, despite the differing illocutionary force. The first conjunct is merely set-up for the second. When these conditions are not met, the coordinations are marked. But it is this markedness that provides the argument against linking coordination and supplementation. The descriptive generalization, the one that should guide the theory, is that we find no such markedness with assertive supplements inside main clauses with nonassertive force.

The fully acceptable discourse in (5.3d) suggests that the problem really does trace back to the nature of coordination. Furthermore, it suggests the falsity of an assumption that is common in dynamic logic, namely, that sequences of sentences not joined by overt *and*s are semantically identical to such sequences joined by *and*. It seems clear that we must locate the semantic and pragmatic differences between, e.g., (5.3c) and (5.3d) in the word *and*.

In static treatments, the difference between (5.3c) and (5.3d) is located in their final denotations. Whereas coordinations of two proposition-denoting expressions denote single propositions, sequences of such sentences denote *tuples* of propositions. Of course, on the CI treatment of supplements, this is true also of sentences that contain clausal supplements. Thus, the $\mathcal{L}_{\mathrm{CI}}$ treatment seems to group these discourse-types properly.

An obvious derivational alternative is to hypothesize that supplement-containing expressions are derived, not from coordinate structures, but rather from sequences of sentences. Ross (1967:§6.2.4.1) suggests that this second analysis might be the correct one, but he is not overly enthused: "it looks like the best analysis of appositives that is presently available, but one which is none too good" (p. 241). It does have the advantage of allowing that supplement-containing sentences denote tuples of meanings. But is it necessary? Chapter 3 shows that we can have all the effects of this multiple-sentence analysis without explicitly deriving the surface strings from abstract underlying sequences of structures, or objects rooted at a text-node, or any of these other transformational alternatives. We can interpret the surface structures. Moreover, the syntactic evidence points to the surface structures as the only ones. Hence, it seems that the sequence-of-sentences approach to supplements, while able to avoid the pitfalls of the coordination analysis, is simply a clumsy restatement of the CI-based analysis.

5.4 The transformational mapping

It seems clear that the right-arrow in the example of McCawley's (1998) analysis in (5.2) hides a considerable amount of complexity. The road between (5.2a) and (5.2b) is surely long. It is clear that it must also vary depending on the kind of supplement we seek to derive.

McCawley (1998:453) specifies that the supplement itself is "moved, without change of constituent structure, to a position immediately following the target" (where 'target' is his term for the anchor). This operation is likely to prove quite computationally demanding, as it involves a global search of the main clause based on semantic coreference. What's more, *As*-parentheticals and niched coordinations show that it is not invariably the case that the supplement must follow its anchor; *As*-parentheticals, for example, may sit either before or after the proposition that saturates their meaning, though the adjacency requirements remain in force, as seen in (5.6) (see Potts 2002b for the details of how this plays out, which depend on the verb-phrase internal subject hypothesis)).

(5.6) a. As the judge wrote, Chuck agreed that the verdict was fair.
b. Chuck agreed that, as the judge wrote, the verdict was fair.
c. Chuck agreed that the verdict was fair, as the judge wrote.

Examples (5.6a) and (5.6b) are each unambiguous. In (5.6a), the interpretation of the *As*-clause must be roughly that the judge wrote that Chuck agreed that the verdict was fair. That is, it must take the entire matrix clause as its argument. In (5.6b), this matrix interpretation is unavailable: the judge wrote that the verdict was fair, but need not have said anything about Chuck at all. Both readings are possible for the clause-final version, (5.6c), in which the string does not tell us whether the *As*-parenthetical is adjoined to the matrix or the embedded sentence.

The situation gets even trickier when we consider niched coordinations such as (5.7).

(5.7) a. Luke has — and this is amazing — eaten fifty eggs.
b. Luke has — and you'll never believe this — eaten fifty eggs.
c. Luke has — and Ed said this blew him away — eaten fifty eggs.

In all cases in (5.7), the deictic element *this* is interpreted as anaphoric to *Luke has eaten fifty eggs*. But its position inside the supplement is highly variable; unlike supplementary relatives and *As*-parentheticals, this construction does not permit a simple search method for finding the relevant coreferential term. Here again, defining the algorithm that McCawley's rightarrow is supposed to stand in for is likely to prove computationally daunting and none too enlightening.

What's more, complexity issues aside, we can already see the essentials of the $\mathcal{L}_{\mathrm{CI}}$ treatment emerging. In the previous section, we saw that we would have to move to a multidimensional treatment; coordination is a poor approximation of the content supplements define. Here, we see that lexical marking is necessary: we will need a host of transformations or sets of transformations, each delivering a different surface structure. These individuated mappings do the work of the type-theoretic distinctions that serve as the foundation for the CI-based treatment. But whereas the type-theoretic approach merely takes advantage of the nature of the usual kind of description logic for linguistic semantics, the transformational view does exactly what transformationalists have always found undesirable: it leans on a set of ad hoc, lexically-conditional transformations. Thus, even without delving into the gory details of the mapping, we can see that the $\mathcal{L}_{\mathrm{CI}}$ alternative is preferable. It does all the positive work of the transformational account, but without any transformations. The function that relates syntactic structures to semantic parsetrees is given in general terms; it hides no lexical marking.

5.5 The surface

Some more recent work on supplements converges on a version of the wide-scope conjunction hypothesis that does not run afoul of coordination arguments, simply because it posits no link with coordination. In addition, at least in the hands of Huddleston and

Pullum (2002), it does away with the transformational mapping completely. Like the $\mathcal{L}_{\mathrm{CI}}$ treatment, this approach seeks to interpret the surface structure. The differences lie in the mechanisms that facilitate this interpretation. This section follows two paths one might take from here. The first, due to McCawley, involves treating supplements (at least those considered here) as entering into dominance relations with the root. The second posits a new supplement relation, one that is disjoint from dominance and restricted so as to involve root-level adjunction only.

The discussion mainly concerns the model theory for the natural language syntax; the central question is, ‘What kind of structures do we need in order to make sense of a supplement relation?’ I seek first to make the question more specific: we really want to ask about the relationship between dominance and precedence. This is a question about how we formulate the structures, so the discussion often concerns different axiomatizations for trees.

5.5.1 Trees

To facilitate discussion, I define a set of structures called simply *trees*. The definition is adapted from Blackburn et al. 2001:6 and is intended as a basis for further articulation.

(5.8) A *tree* is a triple $\mathcal{T} = (T, \geqslant, \prec)$, where,

a. $T = \{u_0, u_1, u_2, \ldots\}$ is a set of nodes.

b. $\geqslant \subseteq T \times T$ is the domination relation. It is reflexive, transitive and antisymmetric.

c. $\gtrdot$ is the immediate domination relation; for all $u, u' \in T, u \gtrdot u'$ iff,

  i. $u \geqslant u'$;

  ii. $u \neq u'$;

  iii. $\forall u''[(u \geqslant u'' \wedge u'' \geqslant u') \rightarrow (u'' \geqslant u \vee u' \geqslant u'')]$

d. Every $u \in T$ has a unique $\gtrdot$-predecessor.

e. T contains a unique r, the root, that is maximal w.r.t. domination: for all $u \in T$, $r \geqslant u$.

f. $\prec \subseteq T \times T$ is the linear succession relation. It is irreflexive, transitive, and asymmetric.

g. $\geqslant$ is disjoint from $\prec^+$ (**exclusivity**).

Not all linguists accept the full set of conditions in (5.8). Sampson (1975) rejects condition (5.8d), the single-mother condition. In a sense, this same move is made in Arc-Pair Grammar (Johnson and Postal 1980), though domination plays a much different role in that framework than it does here. But the above represents the common ground for researchers when it comes to supplements. One way to see to the heart of the issue is to notice that two things are missing from the above definition: (*i*) a condition that says the union of dominance and precedence totally order the set of nodes T (**exhaustiveness**); and (*ii*) a condition ensuring that all precedence relations are inherited down the tree (**nontangling**). The discussion below shows that neither condition can be maintained if supplements enter into dominance relations with the root. But both conditions are maintainable if the trees are enriched with a *supplement relation.*

5.5.2 Supplements and dominance

The following are statements of **nontangling** and **exhaustiveness** (adapted from Partee et al. 1993:441–442; Rogers 1998:16).

(5.9) $\forall w, x, y, z \in T[[w \prec x \wedge w \geqslant y \wedge x \geqslant z] \rightarrow y \prec z]$ **nontangling**

'If w precedes x, then everything that w dominates precedes everything that x dominates.'

(5.10) $\forall x, y \in T[\neg[x \prec y \vee y \prec x] \rightarrow [x \geqslant y \vee y \geqslant x]]$ **exhaustiveness**

'If two nodes are not in any precedence relation with each other, then they are in some dominance relation with each other.'

The pairing of **exhaustiveness** with the **exclusivity** condition (5.8g) ensures that dominance and precedence totally order the set of nodes T. That is, the union of $\geqslant$ and $\prec$ is T^2, where T^2 is the set of all ordered pairs of elements drawn from T. As a result, we can prove the following:

(5.11) **Theorem (Partee et al. 1993:442)** For all $x, y, z \in T$, if $x \gtrdot y$ and $x \gtrdot z$, then $y \prec z$ or $z \prec y$.

I do not repeat the proof here, since it is given in full by Partee et al. (1993). The heart of the argument is that two sister nodes cannot be in a domination relation with each other. By **exhaustiveness**, this implies that they are in some precedence relation with each other.

In conjunction with the **nontangling** condition (5.9), this proof has serious consequences for the surface form in (5.2b), repeated in (5.12).

(5.12)

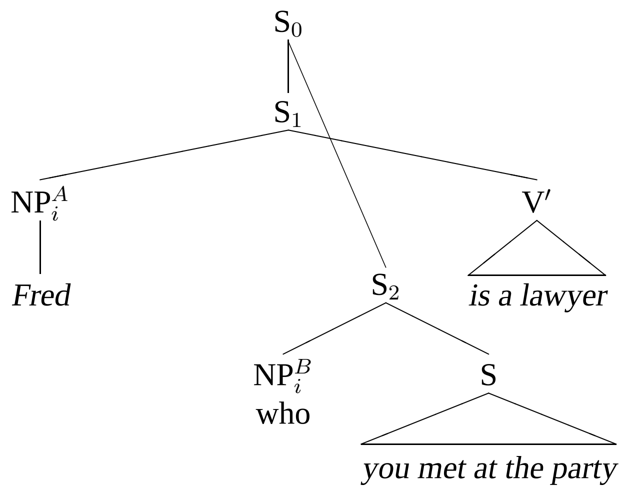

If we interpret the line connecting S_0 with S_2 as a graphic representation of dominance, and we adopt a theory of trees that includes both **exhaustiveness** and **nontangling**, then we have the following oddity: (5.12) is just a misleading picture of the deep structure. It happens that the daughters of the mother node are positioned on the page in an odd and unexpected way in (5.12), but the truth is that the right daughter of the root node, labelled S_2, *follows* the V′ node on this formulation or precedes the NP_i^A node. Here is an argument that brings this out:

(5.13)
i. The daughters of S_0 are S_1 and S_2. By (5.11), one precedes the other.
ii. $S_2 \geqslant NP_i^B$ and $S_1 \geqslant V'$.
iii. Suppose $S_1 \prec S_2$.
iv. Then by **nontangling**, $V' \prec NP_i^B$.

Of course, we should not conclude from this that the transformational mapping represented in (5.2) is an identity. McCawley (1982:93) provides a definition of trees that adheres neither to **nontangling** nor to **exhaustiveness** (see also McCawley 1998:47). The following conditions replace them; the third condition, (5.14c) is of the greatest importance here, since it links dominance and precedence.

(5.14) a. **leaf**$(x) \leftrightarrow$
$$x \in T \wedge \forall y[x \geqslant y \rightarrow x = y]$$
'A node counts as a leaf iff it has no daughters.'

b. $\forall x, y \in T[(\textbf{leaf}(x) \wedge \textbf{leaf}(y)) \rightarrow (x \prec y \vee y \prec x)]$

'The terminals are linearly ordered by precedence.'

c. $$\forall w, x, y, z \in T \left[\begin{pmatrix} \neg\textbf{leaf}(w) \wedge \\ \neg\textbf{leaf}(x) \wedge \\ w \prec x \end{pmatrix} \leftrightarrow \left(\begin{pmatrix} w > y \wedge \\ \textbf{leaf}(y) \wedge \\ x > z \wedge \\ \textbf{leaf}(z) \end{pmatrix} \rightarrow y \prec z \right) \right]$$

'A nonterminal w precedes a nonterminal y iff every terminal that w dominates precedes every terminal that y dominates.'

On McCawley's reworking of nontangling, the S_1 and S_2 in (5.12) do not enter into any precedence relations with each other. Here is a proof of that:

(5.15) a. Assume *party* $\prec$ *is.*

b. Assume *Fred* $\prec$ *who.*

c. Assume $S_1 \geqslant$ *is* and $S_2 \geqslant$ *party*.

d. i. Suppose $S_1 \prec S_2$ (both are $\neg$**leaf**).

ii. Then by (5.14c), *is* $\prec$ *party*. But this contradicts our first assumption.

e. i. Suppose $S_2 \prec S_1$.

ii. Then by (5.14c), *who* $\prec$ *Fred.* But this contradicts the second assumption.

f. Hence, S_1 and S_2 are not ordered w.r.t. $\prec$.

The question before us is whether the correct response is to give up on the usual axioms for trees. McCawley finds motivation for discontinuous structures of the sort allowed by his definition in a variety of constructions: verb–particle structures, relative clause extraposition, right-node raising, and comparatives. However, it seems incorrect to lump supplements with this group, which is itself heterogenous. The next section explores a different strategy: we retain the usual definitions, including **exhaustiveness** and **nontangling**, but add a third primitive binary relation, the supplement relation.

5.5.3 Supplements and the supplement relation

In what follows, I employ the essential insights represented by McCawley's analysis, but I do so by adopting an interpretation of the line connecting S_0 and S_2 as distinct from dominance. Following the terminology of Huddleston and Pullum 2002, I call this the supplement relation. I represent it in structures using dashed lines, and symbolize it with a circled-S, Ⓢ. Thus, a better graphic representation is (5.16).

(5.16)

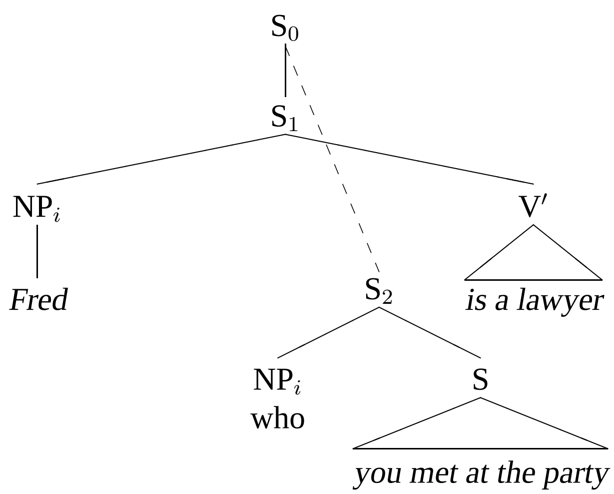

Structures of this form are three-dimensional trees. Traditionally, the horizontal axis represents linear precedence and the vertical axis represents dominance. The geometric z axis, the one that is harder to represent on the page, represents supplementation.

Before proceeding, I want to point out one point of variation between McCawley's (1998) analysis and the one found in Huddleston and Pullum 2002, where the supplement and the root of the main clause are not sisters, as in (5.16), but rather the supplement adjoins directly to the root, as exemplified in (5.17).

(5.17)

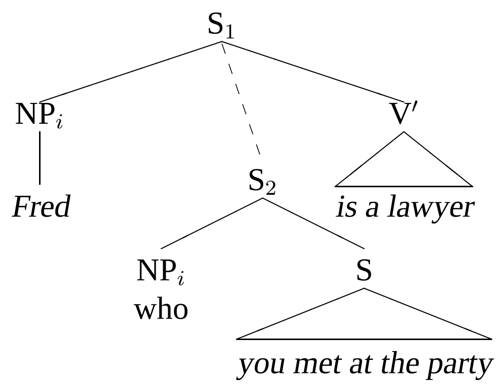

Does it make a difference which structure we work with? Probably not. Consider first what it would take to move from (5.12) to (5.17). We would need a homomorphic mapping that was an identity on every node except S_0 and S_1, which it would map to a single node. All relations would remain the same, except all $\langle \mathrm{S}_0, x\rangle \in (\textcircled{S} \cup \geqslant \cup \prec)$ would go to $\langle \mathrm{S}_1, x\rangle$.

In the direction from (5.17) to (5.16), one might worry that certain command relations would be lost. But this is a remnant of the representation. We must keep sight of the fact that S_2 and V′ (for example) are not sisters in (5.17). What's more, if we did find a reason to talk about the hierarchical relationships between S_2 and V′, we could do so in either structure, using slightly different conditions for each.

The reason this mapping seems harmless is that S_0 is a nonbranching node from the point of view of both dominance and supplementation. Nonbranching syntactic nodes are quite generally eliminable; one would be very hard pressed to find an instance in which a syntactic generalization depended crucially on a node having exactly one daughter.

In what follows, I assume that we could have structures of either form; the conditions on the supplement relation permit either approach.

An essential difference imposed by the supplement relation is that we need to adjust our assumptions about the root. The original condition says that a single node r is maximal w.r.t. domination. In the supplement structures, many nodes might have this status: the root of the main clause, and the root node of any supplement. However, the necessary adjustment to this requirement is minimal. We simply need to allow that the maximal node might be maximal w.r.t. supplement relations:

(5.18) $\forall x \in T[r \geqslant x \vee \exists y[r \circledS y \wedge y \geqslant x]]$ **(connectedness w.r.t.** $\geqslant \cup \circledS$**)**

'The designated root r is such that for every node u, r either dominates u or r bears the supplement relation to a u' that dominates u.'

It is important to ensure also that supplements adjoin only to the root node. The following condition ensures this:

(5.19) $\forall x, y \in T[x \circledS y \rightarrow x = r]$

It is this clause that ensures widest-scope for supplements, in effect demanding that they have the semantic force of a main clause.

Other than this minor modification to the conditions on r, and the addition of an axiom for handling supplements, we need not alter the usual conception of trees in order to include the supplement relation in our theory of syntactic structures. Once again, though, we should ask whether it is appropriate for the constructions addressed here. Given the syntactic integration of all of them, it seems that we should prefer an analysis that sticks to an entirely standard syntax, at least for them. It might be that the supplement relation is appropriate for interjections like *of course* (McCawley 1982), but it seems not to yield the best theory of supplements in general. The next section seeks to reenforce this conclusion based on semantic considerations.

5.5.4 Interpreting supplement structures

Obtaining the proper interpretation from supplement structures is not as straightforward as one might think. One possibility, of course, is to map them to parsetrees of the sort employed throughout this work, letting the workings of $\mathcal{L}_{\mathrm{CI}}$ handle the semantics. But then we are left with the task of formulating this mapping. Like the transformational mapping discussed above, it is likely to be complex. It is also without narrowly syntactic motivation. So the supplement relation would be an unnecessary addition. This seems contrary to the spirit of this syntactic approach, which is to replace the special semantic assumptions of $\mathcal{L}_{\mathrm{CI}}$ with a special syntax. Hence, the question should be, 'How do we do direct model-theoretic interpretation on supplement structures?' It turns out that this question hides genuine complexity as well. Consider the following familiar sort of example:

(5.20) a. Bonnie, who felt jilted, shot Clyde in the head.
b. Bonnie shot Clyde, who felt jilted, in the head.

(5.21) a. That Americans need cheap oil, as Bush observed, is a serious environmental problem.
b. That Americans need cheap oil is a serious environmental problem, as Bush observed.

For each pair, the main clause is the same, yet the supplements are interpreted differently, presumably a result of their linear positioning. Thus, even assuming root attachment, supplements would have to take material they are adjacent to as arguments. The expected method for achieving this result involves adjoining the appositive to the constituent that it receives its interpretation from. The upshot is that adjoining the supplement to the root is by no means the only assumption required to obtain the right meanings. Root adjunction provides only widest-scope (via a stipulation about where supplements can adjoin). It does not tell us what those supplements actually mean. What seems to be required is a semantic rule referencing linear adjacency; the rule in (5.22) does the job. (I use $[\![\cdot]\!]^{\mathcal{M}}$ for a generic interpretation function.)

(5.22) Suppose that the root has all and only the dominance daughters u_d and $u_{d'}$, where $[\![u_{d'}]\!]^{\mathcal{M}}$ is in the domain of $[\![u_d]\!]^{\mathcal{M}}$.
Suppose furthermore that r has supplement daughters $u_{s_1}, \ldots, u_{s_n}$, where each u_{s_i} is immediately adjacent to a node $u_{s_i^*}$.
Then

$$[\![r]\!]^{\mathcal{M}} = \langle [\![u_d]\!]^{\mathcal{M}}([\![u_{d'}]\!]^{\mathcal{M}}), [\![u_{s_1}]\!]^{\mathcal{M}}([\![u_{s_1^*}]\!]^{\mathcal{M}}), \ldots, [\![u_{s_n}]\!]^{\mathcal{M}}([\![u_{s_n^*}]\!]^{\mathcal{M}})\rangle$$

The rule applies the meaning of the root node of a supplement-adjoined subtree to the meaning of its neighbor, and includes the result of this application in the interpretation of the root node. It is now evident that, even having adopted multidimensional syntactic structures, we are still left with a multidimensional semantics. Since supplementation is demonstrably not coordination (see section 5.3), but rather akin to sequences of sentences (a fact $\mathcal{L}_{\mathrm{CI}}$ captures directly) we must have denotations that are the same as those for sequences of sentences, i.e., denotations that are tuples of independent meanings. Moreover, since we need principles to ensure that only a limited class of expressions adjoin via the supplement relation, we have the equivalent of lexical marking as well.

Thus, we duplicate the central tenets of the CI description, but with the addition of a more complex view of the natural language syntax. We seem not to have gained anything. The rule in (5.22) is no more or less suspicious than the interpretive rule that $\mathcal{L}_{\text{CI}}$ relies on, repeated in (5.23).

(5.23) **parsetree interpretation**

Let $\mathcal{T}$ be a semantic parsetree with the at-issue term $\alpha : \sigma^a$ on its root node, and distinct terms $\beta_1 : \langle s^a, t^c \rangle, \ \ldots, \ \beta_n : \langle s^a, t^c \rangle$ on nodes in it (extensionally, $\beta_1 : t^c, \ \ldots, \ \beta_n : t^c$). Then the interpretation of $\mathcal{T}$ is the tuple

$$\left\langle [\![\alpha : \sigma^a]\!]^{\mathcal{M}_i,g}, [\![\beta_1 : \langle s^a, t^c \rangle]\!]^{\mathcal{M}_i,g}, \ldots, [\![\beta_n : \langle s^a, t^c \rangle]\!]^{\mathcal{M}_i,g} \right\rangle$$

where $[\![\cdot]\!]^{\mathcal{M}_i,g}$ is the interpretation function, taking formulae of the meaning language to the interpreted structure $\mathcal{M}_i$, relative to a variable assignment g.

What's more, it seems a mistake to try to assess the relative complexity of (5.22) relative to its nearest $\mathcal{L}_{\text{CI}}$ counterpart, the rule of **CI application**:

(5.24) **CI application**

$$\begin{array}{ccc}
 & \beta : \sigma^a & \\
 & \bullet & \\
 & \alpha(\beta) : \tau^c & \\
 \diagup & & \diagdown \\
\alpha : \langle \sigma^a, \tau^c \rangle & & \beta : \sigma^a \\
\bullet & & \bullet \\
\gamma : \rho^c & & \delta : \upsilon^c
\end{array}$$

From some perspectives, (5.24) is preferable in that it does not reference linear order. But if we were working in a framework like categorial grammar, where linear order is crucial for interpretation, the reverse might be the case, with (5.23) appearing suspicious when set alongside (5.22).

But since the semantics seems entirely comparable for each system, we should ask what work the supplement relation is doing for us. It complicates the syntax, and it will demand a restatement of the syntactic facts uncovered and explored in section 3.4, all of which points to a conservative, adjunction-based syntax for supplements.

5.5.5 In sum

One senses that this analysis does not boast much in the way of explanatory force. We will require lexical marking or some related notion to distinguish supplement structures from coordination, and we will anyway need to move to a multidimensional theory to capture the independence of supplementary content from regular content. That is, we need to duplicate the central notions of the CI logic. It quickly becomes apparent that wide-scope conjunction, while capturing some valid intuitions, is a clumsy restatement of the CI logic's guiding ideas.

Thus, once again, we see the essentials of $\mathcal{L}_{\mathrm{CI}}$ emerging in this syntactic setting. A proposal based on the supplement relation is not at odds with the CI hypothesis (though, with some bluster, one could frame it as such). Rather, the supplement relation locates the distinction between regular and CI content in the syntax, rather than the semantics. Descriptively, there is no doubt that supplements have special characteristics. The question is what the nature of these characteristics is. Here is a slogan: we require a multidimensionality. But at which level of grammatical analysis? The syntax or the semantics? This is not merely a border dispute. The answers have diverse and apparent factual consequences.

Chapter 6

A look outside Grice’s definition

6.1 Neighboring territory

To conclude this dissertation, I think it is useful to step outside of the definition of conventional implicatures (CIs). This move yields new insights into the range of possible natural language meanings and how they relate to Grice’s definition. It also sheds new light on CIs; by inspecting them from the outside, as it were, we are able to see them afresh.

So this chapter takes a look at what happens when we remove individual clauses from Grice’s (1975) definition of CIs, which I repeat in (6.1).

(6.1)
a. CIs are part of the conventional (lexical) meaning of words.
b. CIs are commitments, and thus give rise to entailments.
c. These commitments are made by *the speaker of the utterance* “by virtue of the meaning of” the words he chooses.
d. CIs are logically and compositionally independent of what is “*said* (in the favored sense)”, i.e., independent of the at-issue entailments.

I look closely only at all the possible three-clause versions of this definition, subtracting each property one at a time and trying to determine what, if anything, results.

6.2 Minus lexicality

If we remove the lexicality property (6.1a), we find ourselves in a world of commitments (logical entailments) that inhere in no specific lexical item. I am uncertain that this is, or could be, a genuine class of meanings. I conjecture that the closest that natural languages come to such meanings are things like sarcasm, skepticism, and the like. But these are quite generally signalled by a specific intonation contour. Throughout chapter 3, I rely on the notion that specific intonation contours can give rise to specific meanings. I argue that these intonation contours thus belong in the lexicon. Precedents for this decision include the vast literature on focus semantics and Gunlogson's (2001) recent work on rising declaratives (*Ed dates a gorilla?*). Thus, we would require an overly restrictive notion of what it means to be 'lexical' to be able to count sarcasm among the nonlexical, speaker-oriented, multidimensional entailments.

I nonetheless regard this as a useful experiment. It suggests that when we cannot leave the realm of lexical meanings without also entering into the fuzzier world of noncommitments. Nonnegotiable meanings, it seems, must grounded in the lexicon.

6.3 Minus commitment

When we remove (6.1b), we step into the world of pragmatics, where the primary kind of meaning is the conversational implicature. Of course, as Grice (1975) originally defined conversational implicatures, they also lack the lexicality property (6.1a). If this is the correct way to go, then we can combine this conclusion with that of the previous section to form the generalization that a meaning is lexical if, and only if, it is an entailment.

Chierchia's (2001) neo-Gricean view of conversational implicatures, discussed in chapter 1, section 1.4.1, calls into question the validity of the implication 'if a meaning is lexical, then it is a commitment'. Chierchia builds scalar conversational implicatures into the lexicon in the following limited sense: a connective like *or* has a conversational implicature dimension that is equivalent to *not and*. This permits one to derive scalar implicatures in tandem with the regular compositional at-issue semantics. However, as I discussed in chapter 1, this is not a replacement for Gricean pragmatic theory. Chierchia provides a calculus for deriving some potential meanings, but we must still call upon principles of cooperative social interaction to determine whether these potential meanings become actual.

Removing the commitment property from the definition of CIs helps to highlight a way in which CIs are particularly amenable to a formal treatment. Because they are commitments, we are able to avoid many of the discourse-related considerations that make it so hard to submit conversational implicatures to the usual mode of theorizing in formal semantics. When we talk about CIs, we talk about the semantics proper, and can thus avail ourselves of the major innovations of that subfield of linguistics.

6.4 Minus speaker-orientation

It might strike readers as surprising that this dissertation has so far had little to say about the textbook examples of CIs, which I highlight in (6.2).

(6.2)
a. Shaq is huge <u>but</u> he is agile. (Bach 1999:(1))
b. He is an Englishman; he is, <u>therefore</u>, brave. (Grice 1975:44)
c. <u>Even</u> George$_{\mathbf{F}}$ could pass the test.
d. I am still not happy with the examples.
e. Adam attacked again.

These examples have been relegated to minor roles in this work because when we test them for the properties in Grice's (1975) definition, repeated in (6.1), we find that (6.1c) is missing. It seems likely that Grice overlooked this property because he almost never investigated embedded examples. In general, he used only monoclausal utterances, which tend to make everything speaker-oriented simply because they do not supply any other agent to whom the content might be attributed. This oversight is not limited to Grice's study of CIs. Chierchia (2001) notes that it plagues his theory of conversational implicatures as well: Grice's writing does not recognize that conversational implicatures can arise from embedded constituents, making them relations among proper subparts of utterances.

When we take away clause (6.1c), we are left with meanings that are multidimensional but for which we find little motivation, factual or conceptual, for leaving the realm of at-issue content. Let's call such expressions *merely multidimensional*. Examples with *but* are classics, so let's inspect the meaning of example (6.2a) more closely. In chapter 1, section 1.4.5, I adapted and expanded the insights of Bach (1999) to show that *but* is, in present terms, merely multidimensional, as suggested by (6.3).

(6.3) a. Shaq is huge <u>but</u> he is agile. (Bach 1999:(1))

b. primary entailment: $\textbf{huge}(\textbf{shaq}) \wedge \textbf{agile}(\textbf{shaq})$

c. ancillary entailment: $\mathsf{G}x[\textbf{huge}(x) \rightarrow \neg\textbf{agile}(x)]$

I propose that we use *primary entailment* to pick out the most prominent entailment of any sentence, and *ancillary entailment* for all nonprimary entailments (with *secondary*, *tertiary*, and the like working in the expected way). I will leave open whether it is useful to classify CIs as ancillary entailments. The logic keeps CIs distinguished from ancillary entailments in the grammar, but when we get to the models, we could define projection functions that pick out, say, any noninitial member in a tuple of meanings. At that point, ancillary entailments and CIs could have the same status.

As I discussed in chapter 1, what I have here called the ancillary entailment for *but* is contextually variable. Sometimes, the nature of the sentence and the default assumptions about the context make the generic quantification of (6.3c) highly implausible:

(6.4) a. He's from New York, but from Upstate.

b. No students went to the party, but many professors did.

The generic quantification seems to be the one that arises in basic cases. So I adopt it as the proper secondary meaning for *but.* The important point is not the nature of the proposition expressed, but rather than there is a propositional meaning there.

Bach (1999) makes a central observation: both contributions of *but* are embeddable under propositional attitude predicates, particularly those that establish indirect discourse contexts. The following examples are from Bach 1999:

(6.5) a. Marv said that Shaq is huge *but* that he is agile.

b. Marv believes that being huge is a good indicator of agility. #Marv said that Shaq is huge *but* that he is agile.

This points up a significant contrast with the CI expressions analyzed in the preceding chapters, which do not embed even under *say*. The nominal appositive in (6.6) illustrates:

(6.6) a. Sheila believes that Chuck, a confirmed psychopath, is fit to watch the kids.

b. = Chuck is a confirmed psychopath, and Sheila believes that Chuck is fit to watch the kids.

c. ≠ Sheila believes that Chuck is a confirmed psychopath and that Chuck is fit to watch the kids.

What this indicates is that we must allow verbs like *say* to target both dimensions of the meaning of *but*. But we must find a way to do this that respects the independence of the primary meaning from any ancillary meanings, and any ancillary meanings from each other. Analyses in earlier chapters employ just such a meaning for *say*. In (6.7), I repeat the generalized, intensional translation that I provided in chapter 2.

(6.7) $$\textit{say} \rightsquigarrow \lambda\langle \mathbf{p}_1, \ldots, \mathbf{p}_n\rangle \lambda x. \left\langle \begin{array}{c} \lambda w.\, \mathbf{say}_w(\mathbf{p}_1)(x), \\ \vdots \\ \lambda w.\, \mathbf{say}_w(\mathbf{p}_n)(x) \end{array} \right\rangle : \langle \overrightarrow{\langle s^a, t^a\rangle}_n, \langle e^a, \overrightarrow{\langle s^a, t^a\rangle}_n\rangle\rangle$$

Recall that $\overrightarrow{\langle s^a, t^a\rangle}_n$ abbreviates a product type consisting of two propositional types: $\langle\langle s^a, t^a\rangle, \langle s^a, t^a\rangle\rangle$. The types in the meaning for *say* are all at-issue. But the logic $\mathcal{L}_{\mathrm{CI}}$ yields a multidimensional meaning for clauses in which it appears. I illustrate in (6.8), which employs only extensional types for the sake of simplicity.

(6.8) Marv said that Shaq is huge *but* that he is agile.

$$\frac{\textbf{marv} : e^a \qquad \dfrac{\textbf{say} : \langle\langle t^a \times t^a\rangle, \langle e, \langle t^a \times t^a\rangle\rangle\rangle \qquad \textbf{but}(\textbf{agile}(\textbf{shaw}))(\textbf{huge}(\textbf{shaq})) : \langle t^a \times t^a\rangle}{\textbf{say}(\textbf{but}(\textbf{agile}(\textbf{shaw}))(\textbf{huge}(\textbf{shaq}))) : \langle e, \langle t^a \times t^a\rangle\rangle}}{\textbf{say}(\textbf{but}(\textbf{agile}(\textbf{shaw}))(\textbf{huge}(\textbf{shaq})))(\textbf{marv}) : t^a}$$

By the rule for interpreting trees, **parsetree interpretation**, this denotes $\langle 1, 1\rangle$ just in case Marv said (*i*) that Shaq is huge and agile; and (*ii*) that hugeness generally precludes agility. This is the result that Bach (1999) calls for.

Not all propositional attitude verbs have the property that they take tuples of meanings as arguments. Verbs of wagering provide a nice counterpoint to *say*. They seem to target only the primary assertion, as illustrated in (6.9).

(6.9) Ed bets that Thora is small but loud.

With *Thora is small but loud*, Ed suggests that smallness generally precludes loudness. Inexperienced with small babies, Ed does not realize that this is false: small humans are generally loud humans. Does Ed, therefore, lose this bet? Most speakers say that only the question of whether Thora is at the intersection of the small things and the loud things impacts this question. The secondary entailment is not relevant; the content of the bet concerns only the primary at-issue entailment. Again, using the product types of $\mathcal{L}_{\mathrm{CI}}$, we can formulate meanings of the requisite sort:

(6.10) *bet* $\rightsquigarrow$

$$\lambda\langle \mathbf{p}_1, \ldots, \mathbf{p}_n\rangle \lambda x. \left\langle \begin{array}{c} \lambda w.\, \textbf{bet}_w(\mathbf{p})(x), \\ \vdots \\ \mathbf{p}_n \end{array} \right\rangle : \langle \overrightarrow{\langle s^a, t^a\rangle}_n, \langle e, \overrightarrow{\langle s^a, t^a\rangle}_n\rangle\rangle$$

Francescotti (1995) makes exactly this observation for the behavior of *even*. Of (6.11a, b), he writes, "Suppose that we are placing bets on whether Albert failed the exam. Feeling confident that he did fail, I utter sentence (6.11a). Suppose, however, that Albert's failing is not at all surprising, and in fact is very likely. In this case, (6.11a) would certainly be inappropriate. However, assuming that Albert did fail, it seems odd to think that (6.11a) is false, and that I should therefore pay up" (p. 153).

(6.11) a. Even Albert failed the exam.
b. Albert failed the exam.

Much of Francescotti's (1995) attention is given over to the question of whether *even* contributes to the "truth conditions". As is well-known (and as Francescotti reviews in much detail), this question has been asked and answered many times before, never conclusively. Such situations rarely resolve in favor of one position or the other. Usually, the dispute itself is shown to have begun with a false premise. I argue that this is the case for *even*. The false premise is that the "truth-conditions" must be one-dimensional — a single proposition, say. From the present perspective, (6.11a) can be assigned a multidimensional meaning. The primary entailment is equivalent to (6.11b). The secondary entailment is notoriously slippery in this case. But the exact nature of that content is not pressing. The point is that we can have analyses such as (6.12) (which provides a meaning for *even* as an adnominal modifier only, to keep things simple).

(6.12) a. *even* $\rightsquigarrow \lambda x \lambda f. \left\langle \begin{array}{c} f(x), \\ \textbf{among-the-least-likely-to-}f(x) \end{array} \right\rangle$

b. $\left\langle \begin{array}{c} \textbf{pass}(\textbf{the}(\textbf{test}))(\textbf{albert}), \\ \textbf{among-the-least-likely-to-pass}(\textbf{the}(\textbf{test}))(\textbf{albert}) \end{array} \right\rangle :$
$\langle t^a \times t^a \rangle$

- $\textbf{even}(\textbf{albert}) : \langle\langle e^a, t^a\rangle, \langle t^a \times t^a\rangle\rangle$
  - $\textbf{even} : \langle e^a, \langle\langle e^a, t^a\rangle, \langle t^a \times t^a\rangle\rangle\rangle$
  - $\textbf{albert} : e^a$
- $\textbf{pass}(\textbf{the}(\textbf{test})) : \langle e^a, t^a\rangle$

The meaning for *bet* given in (6.10) now takes over to derive Francescotti's observations about (6.11a) when money is on the line. The *bet* defined there can take this product type as an argument, but, in an intuitive sense, it applies only to the meaning of the first coordinate of the pair of meanings.

A great many constructions count as merely multidimensional. Other likely candidates are the additive modifiers (*too, also, as well*, and their negative counterpart *either*), exceptive constructions (e.g., *No Muppet but Kermit*; Hoeksema 1995; von Fintel 1993; Moltmann 1995), and perhaps the multidimensionality of the definite article, as observed by Hawkins (1991) and von Fintel (2001). Detailed study of these cases and their implications for the nature of the at-issue dimension must wait for another occasion.

6.5 Minus multidimensionality

Where we end up after removing (6.1d) is not fully determined. If we treat the removal of (6.1d) as merely the removal of the *independence* of the dimensions in question, then we end up talking about speaker-presuppositions — i.e., those presuppositions that seem not to embed. In this area, the non-at-issue dimension of meaning places strict conditions on the nature of the at-issue dimension that it is associated with. In general, the presuppositions must be true for the expression to have a defined value at all. We could, at least in principle, stick with a multidimensional logic, but we would require many special conditions on the range of possible values for those dimensions. This is, in essence, the move that Karttunen and Peters (1979) make.

But we could also regard the removal of this clause as a return to the view that sentence meanings are always single propositions — the 'one sentence, one proposition' view of Bach (1999). The subclass of such meanings that also possess the other properties specified in (6.1) are the main clause, at-issue assertions. However, both CIs and the merely multidimensional meanings discussed in the previous section show that one-dimensionality is a fiction. Natural language sentence meanings can be much more complex than this.

Nonetheless, it can be a useful fiction. Very often, it is fruitful to pretend that a sentence's value is simply the interpretation of the (at-issue) term on the root node of its parsetreee. For this reason, $\mathcal{L}_{\mathrm{CI}}$ is defined in such a way that the at-issue and CI dimensions can be fairly easily teased apart. The most basic reflection of such a division is in the type definition, where it is possible to isolate the at-issue type specifications from the CI type specifications. If we cut the definition down by two clauses (clauses (ii) and (iv)), we arrive at a standard type definition for at-issue semantics. Similarly, when we look at the set of tree-admissibility conditions, we see that only **CI application** and **feature semantics** can introduce CI content. If we remove the first and amend the second, we have a theory of at-issue meanings.

6.6 In sum

The above discussion amounts to a check of the meanings that are logically and conceptually near to CIs. It might be that the lexicality and commitment properties are not independent of each other. Multidimensionality and speaker-orientation certainly are quite different properties, though. When we remove speaker-orientation, we end up with a class of meanings that $\mathcal{L}_{\mathrm{CI}}$ is capable of describing without difficulty. Removing multidimensionality is probably a fiction, but a useful one. A system for dealing with such meanings is embedded inside $\mathcal{L}_{\mathrm{CI}}$. If we remove the CI and product types, we are left with the means for deriving phrases with single meanings as their denotations. But the preceding chapters clearly indicate that, at the intersection of the multidimensional, the lexical, and the speaker-oriented, we find meanings that are widely attested and ripe for exploration. For such exploration, we'll require all the tools that $\mathcal{L}_{\mathrm{CI}}$ provides.

Appendix A

The logics $\mathcal{L}_{\mathbf{CI}}$ and $\mathcal{L}_U$

A.1 Overview

The primary descriptive tool of this dissertation is $\mathcal{L}_{\text{CI}}$, a higher-order lambda calculus with at-issue and conventional implicature (CI) types (section A.2.1.1). We can regard $\mathcal{L}_{\text{CI}}$ as the specification of a class of semantic parsetrees (section A.2.2). We interpret the terms of $\mathcal{L}_{\text{CI}}$ in intensional models (section A.2.4) using the interpretation function $[\![\cdot]\!]^{\mathcal{M}_i,g}$ (section A.2.5), which is always relativized to some intensional model $\mathcal{M}_i$ (drawn from a specific set $\mathfrak{M}$ of such models) and an assignment g.

The dissertation also calls upon the logic $\mathcal{L}_U$ (section A.3). It is a language with types for declaratives and interrogatives, as well as types for entities and truth values (section A.3.1.1). We interpret $\mathcal{L}_U$ in discourse structures, which model discourses. These models contain a function that maps from discourse participants to intensional models of the sort designed for interpreting $\mathcal{L}_{\text{CI}}$.

The terms of $\mathcal{L}_U$ are distinguished from the terms of $\mathcal{L}_{\text{CI}}$. All $\mathcal{L}_U$ terms are given with raised corner brackets around them, whereas $\mathcal{L}_{\text{CI}}$ terms are bare. But the logic $\mathcal{L}_U$ and its models are connected to the logic $\mathcal{L}_{\text{CI}}$ and its models in three central ways:

(A.1)
a. The discourse structures for $\mathcal{L}_U$ contain the intensional models for $\mathcal{L}_{\text{CI}}$.
b. The discourse structures for $\mathcal{L}_U$ contain sentences, where sentences are pairs of trees, one of which is a semantic parsetree for $\mathcal{L}_{\text{CI}}$. In essence, certain constants of $\mathcal{L}_U$ are interpreted as the parsetrees for $\mathcal{L}_{\text{CI}}$.
c. The interpretation function $[\![\cdot]\!]^{\mathcal{D},s,a}$ for a discourse structure is defined for terms of $\mathcal{L}_U$ as well as entire parsetrees of $\mathcal{L}_{\text{CI}}$. Thus, it functions to connect the two logics at the level of interpretation (section A.4).

It is helpful to think of $\mathcal{L}_U$ as the upper layer in a two-tiered logic, with $\mathcal{L}_{\text{CI}}$ serving as the lower layer.

A.2 The logic $\mathcal{L}_{\text{CI}}$

A.2.1 The syntax of $\mathcal{L}_{\text{CI}}$

A.2.1.1 Types for $\mathcal{L}_{\text{CI}}$

(A.2)
a. e^a, t^a, and s^a are basic at-issue types for $\mathcal{L}_{\text{CI}}$.
b. e^c, t^c, and s^c are basic CI types for $\mathcal{L}_{\text{CI}}$.
c. If σ and τ are at-issue types for $\mathcal{L}_{\text{CI}}$, then $\langle\sigma,\tau\rangle$ is an at-issue type for $\mathcal{L}_{\text{CI}}$.
d. If σ is an at-issue type for $\mathcal{L}_{\text{CI}}$ and τ is a CI type for $\mathcal{L}_{\text{CI}}$, then $\langle\sigma,\tau\rangle$ is a CI type for $\mathcal{L}_{\text{CI}}$.
e. If σ and τ are at-issue types for $\mathcal{L}_{\text{CI}}$, then $\langle\sigma \times \tau\rangle$ is a product type for $\mathcal{L}_{\text{CI}}$, a subset of the set of at-issue types for $\mathcal{L}_{\text{CI}}$.
f. The full set of types for $\mathcal{L}_{\text{CI}}$ is the union of the at-issue and CI types for $\mathcal{L}_{\text{CI}}$.

(A.3) Let x serve as a variable over $\{e, t, s\}$, and let σ and τ serve as variables over well-formed types with their superscripts stripped off. The type-superscript abbreviator $\leftrightsquigarrow$ is defined as follows:

$$\begin{array}{rcl} x^a & \leftrightsquigarrow & x^a \\ x^c & \leftrightsquigarrow & x^c \\ \langle \sigma^a, \tau^a \rangle & \leftrightsquigarrow & \langle \sigma, \tau \rangle^a \\ \langle \sigma^a, \tau^c \rangle & \leftrightsquigarrow & \langle \sigma, \tau \rangle^c \end{array}$$

A.2.1.2 Terms for $\mathcal{L}_{\text{CI}}$

(A.4) Let ME_τ denote the set of all meaningful expressions of type τ for $\mathcal{L}_{\text{CI}}$.

i. If c is a constant of type τ, then $c \in ME_\tau$.

ii. If x is a variable of type τ, then $x \in ME_\tau$.

iii. If $\alpha \in ME_{\langle \sigma^a, \tau^a \rangle}$ and $\beta \in ME_{\sigma^a}$, then $(\alpha(\beta)) \in ME_{\tau^a}$.
If $\alpha \in ME_{\langle \sigma^a, \tau^c \rangle}$ and $\beta \in ME_{\sigma^a}$, then $(\alpha(\beta)) \in ME_{\tau^c}$.

iv. If $\alpha \in ME_{\tau^a}$ and x is a variable in ME_{σ^a}, then $(\lambda x.\, \alpha) \in ME_{\langle \sigma^a, \tau^a \rangle}$.
If $\alpha \in ME_{\tau^c}$ and x is a variable in ME_{σ^a}, then $(\lambda x.\, \alpha) \in ME_{\langle \sigma^a, \tau^c \rangle}$.

v. If $\alpha \in ME_{\sigma^a}$ and $\beta \in ME_{\tau^a}$, then $\langle \alpha, \beta \rangle \in ME_{\sigma^a \times \tau^a}$.

vi. If $\alpha, \beta \in t^a$, then $\neg\alpha$, $(\alpha \vee \beta) \in ME_{t^a}$.
If $\alpha, \beta \in t^c$, then $\neg\alpha$, $(\alpha \vee \beta) \in ME_{t^c}$.

vii. If $\alpha \in ME_{t^a}$, and x is a variable, then $(\exists x[\alpha]), (\mathsf{G}x[\alpha]) \in ME_{t^a}$.
If $\alpha \in ME_{t^c}$, and x is a variable, then $(\exists x[\alpha]), (\mathsf{G}x[\alpha]) \in ME_{t^c}$.

viii. The full set ME of meaningful expressions for $\mathcal{L}_{\text{CI}}$ is the union of all the sets ME_τ for all types τ.

To simplify the definition, I leave out the connectives $\wedge$, $\rightarrow$, $\leftrightarrow$, the quantifier $\forall$, and the definite-description operator $\imath$. These are definable in the expected way in terms of the above clauses.

A.2.1.3 Variable conventions for $\mathcal{L}_{\mathrm{CI}}$

(A.5)

type	variables	name
s	$\{w, w', w'', \ldots\}$	worlds
e	$\{x, y, z, \ldots\}$	individuals
t	$\{p, q, \ldots\}$	truth values
$\langle t, t\rangle$	$\{P, Q, \ldots\}$	sets of truth values
$\langle s, t\rangle$	$\{\mathbf{p}, \mathbf{q}, \ldots\}$	propositions
$\langle\langle s, t\rangle, t\rangle$	$\{\mathbf{P}, \mathbf{Q}, \ldots\}$	set of propositions
$\langle\langle s, t\rangle, \langle s, t\rangle\rangle$	$\{\mathcal{P}, \mathcal{Q}, \ldots\}$	relations on propositions
$\langle e, t\rangle$	$\{f, g, h, \ldots\}$	one-place functions on entities
$\langle s, \langle e, t\rangle\rangle$	$\{\mathbf{f}, \mathbf{g}, \mathbf{h}, \ldots\}$	properties
$\langle\langle e, t\rangle, t\rangle$	$\{F, G, H, \ldots\}$	generalized quantifiers (GQs)
$\langle\langle s, \langle e, t\rangle\rangle, t\rangle$	$\{\mathbf{F}, \mathbf{G}, \mathbf{H}, \ldots\}$	intensional GQs
$\langle s, \langle e, t\rangle, \langle e, t\rangle\rangle$	$\{\mathcal{F}, \mathcal{G}, \mathcal{H}, \ldots\}$	relations on properties
$\langle e, \langle e, t\rangle\rangle$	$\{R, S, T, \ldots\}$	two-place relations
$\langle s, \langle e, \langle e, t\rangle\rangle\rangle$	$\{\mathbf{R}, \mathbf{S}, \mathbf{T}, \ldots\}$	intensional two-place relations

To help prevent orthographic overload, world variables are given as subscripts on the names of functions. Thus:

(A.6) $\lambda x \lambda w.\, \textbf{whine}(x)(w) \leftrightsquigarrow \lambda x \lambda w.\, \textbf{whine}_w(x)$

This names a function that takes any individual x to the set of worlds in which x whines.

I almost always drop outermost parentheses. I assume also that application associates to the left. Thus, $((\alpha(\beta))(\gamma))$ abbreviates to $(\alpha(\beta))(\gamma)$ by the convention that drops outermost parentheses. We can abbreviate further to $\alpha(\beta)(\gamma)$ by the convention that associates application to the left.

A.2.1.4 Abbreviated lexicon of constants for $\mathcal{L}_{\text{CI}}$

The wavy arrow, $\rightsquigarrow$, is the translation function, taking natural language expressions to $\mathcal{L}_{\text{CI}}$.

A.2.1.4.1 Individual constants

(A.7) a. *Ed* $\rightsquigarrow$ **ed** : e^a (*representative individual constant*)

b. *flee* $\rightsquigarrow$ **flee** : $\langle e^a, t^a \rangle$ (*representative predicate*)

c. *write* $\rightsquigarrow$ **write** : $\langle e^a, \langle e^a, t^a \rangle\rangle$ (*representative two-place relation*)

A.2.1.4.2 Propositional-attitude verbs

(A.8) a. *say* $\rightsquigarrow$

$$\lambda\langle \mathbf{p}_1, \ldots, \mathbf{p}_n\rangle \lambda x. \left\langle \begin{array}{c} \lambda w.\, \textbf{say}_w(\mathbf{p}_1)(x), \\ \vdots \\ \lambda w.\, \textbf{say}_w(\mathbf{p}_n)(x) \end{array} \right\rangle : \langle \overrightarrow{\langle s, t\rangle}_n, \langle e, \overrightarrow{\langle s, t\rangle}_n \rangle\rangle$$

(2.27)

b. *bet* $\rightsquigarrow$

$$\lambda\langle \mathbf{p}_1, \ldots, \mathbf{p}_n\rangle \lambda x. \left\langle \begin{array}{c} \lambda w.\, \textbf{bet}_w(\mathbf{p})(x), \\ \vdots \\ \mathbf{p}_n \end{array} \right\rangle : \langle \overrightarrow{\langle s^a, t^a\rangle}_n, \langle e, \overrightarrow{\langle s^a, t^a\rangle}_n \rangle\rangle$$

(6.10)

A.2.1.4.3 Feature terms

(A.9) a. $^{\cap} \stackrel{\textit{def}}{=} \lambda X.\, \imath x[\forall y[X(y) \leftrightarrow y \leqslant x]] : \langle\langle \sigma^a, t^a\rangle, e^a\rangle$ (3.26a)

b. $^{\cup} \stackrel{\textit{def}}{=} \lambda x \lambda Y.\, Y \leqslant x : \langle e^a, \langle \sigma^a, t^a\rangle\rangle$ (3.26b)

c. COMMA $\rightsquigarrow \lambda X \lambda x.\, X(x) : \langle\langle \sigma^a, t^a\rangle, \langle \sigma^a, t^c\rangle\rangle$, where $\sigma \in \{e, t\}$ (3.114)

d. **lift** $\stackrel{\textit{def}}{=} \lambda x \lambda f.\, f(x)$ (4.43)

A.2.1.4.4 Adjectives and common nouns

(A.10) a. $\left\{\begin{array}{c}\textit{damn}\\ \textit{bloody}\\ \vdots\\ \textit{fucking}\end{array}\right\} \rightsquigarrow \lambda X.\, \textbf{bad}({}^{\cap}X) : \langle\langle\tau^a, t^a\rangle, t^c\rangle$ (4.35)

b. $\textbf{bastard} : \langle e^a, t^c\rangle$ (4.45)

A.2.1.4.5 Adverbs

(A.11) a. *luckily* $\rightsquigarrow \left\{\begin{array}{l}\lambda f\lambda x.\, \textbf{lucky}(f(x)) : \langle\langle e^a, t^a\rangle, \langle e^a, t^a\rangle\rangle\\ \lambda p.\, \textbf{lucky}(p) : \langle t^a, t^a\rangle\end{array}\right.$ (3.124)

b. $\lambda \mathbf{p}.\, \textbf{thoughtfully}(\mathbf{p})(x_1) : \langle\langle s^a, t^a\rangle, \langle s^a, t^a\rangle\rangle$

(where x_1 denotes a salient individual; (3.136))

A.2.1.4.6 Others

(A.12) a. $\textbf{kI} \overset{\mathit{def}}{=} \lambda\mathbf{p}\lambda w.\, \exists w'[w\mathsf{E}w' \wedge \neg\mathbf{p}(w')] : \langle\langle s^a, t^a\rangle, \langle s^a, t^c\rangle\rangle$ (4.81)

b. *as* $\rightsquigarrow$

$$\left\{\begin{array}{l}\textbf{as}_{\mathrm{CP}} \overset{\mathit{def}}{=} \lambda\mathcal{P}\lambda\mathbf{p}\lambda w.\, \mathbf{P}(\mathbf{p})(w) :\\ \qquad \langle\langle\langle s^a, t^a\rangle, \langle s^a, t^a\rangle\rangle, \langle\langle s^a, t^a\rangle, \langle s^a, t^a\rangle\rangle\rangle\\ \textbf{as}_{\textbf{manner}} \overset{\mathit{def}}{=} \lambda p\lambda f\lambda x.\, \exists P[P(p) \wedge P(f(x))] :\\ \qquad \langle t^a, \langle\langle e^a, t^a\rangle, \langle e^a, t^a\rangle\rangle\rangle\end{array}\right.$$

(3.112)

c. *but* $\rightsquigarrow$

$$\lambda X\lambda Y\lambda x.\left\langle\begin{array}{c}X(x) \wedge Y(x),\\ \mathsf{G}y[Y(y) \rightarrow \neg X(y)]\end{array}\right\rangle :$$

$$\langle\langle\tau^a, t^a\rangle, \langle\langle\tau^a, t^a\rangle, \langle\tau^a, \langle t^a \times t^a\rangle\rangle\rangle\rangle$$

A.2.1.5 Tree-admissibility conditions for $\mathcal{L}_{\mathbf{CI}}$

The lowercase Greek letters α, β, γ, and δ range over terms. The lowercase Greek letters ρ, σ, τ, and υ range over types.

The bullet, $\bullet$, is a metalogical device for separating independent terms of $\mathcal{L}_{\text{CI}}$. It has no interpretation.

Optional material is inside dotted lines. Hence, all the rules except the first abbreviate a set of rules. We have to allow that there might be CI content hanging around.

(A.13) $\alpha : \sigma$ (where α is a meaningful expression of $\mathcal{L}_{\text{CI}}$)

(A.14) **at-issue application**

$$\begin{array}{ccc} & \alpha(\beta) : \tau^a & \\ \alpha : \langle \sigma^a, \tau^a \rangle & & \beta : \sigma^a \\ \bullet & & \bullet \\ \boxed{\gamma : \rho^c} & & \boxed{\delta : \upsilon^c} \end{array}$$

(A.15) **at-issue intersection**

$$\begin{array}{ccc} & \lambda X.\, \alpha(X) \wedge \beta(X) : \langle \sigma^a, t^a \rangle & \\ \alpha : \langle \sigma^a, t^a \rangle & & \beta : \langle \sigma^a, t^a \rangle \\ \bullet & & \bullet \\ \boxed{\gamma : \rho^c} & & \boxed{\delta : \upsilon^c} \end{array}$$

(A.16) **CI application**

$$\begin{array}{ccc} & \beta : \sigma^a & \\ & \bullet & \\ & \alpha(\beta) : \tau^c & \\ \alpha : \langle \sigma^a, \tau^c \rangle & & \beta : \sigma^a \\ \bullet & & \bullet \\ \boxed{\gamma : \rho^c} & & \boxed{\delta : \upsilon^c} \end{array}$$

(A.17) **isolated CIs**

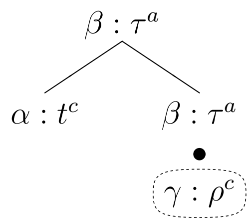

(A.18) **feature semantics**

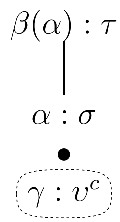

(where β is a designated *feature term* of type $\langle \sigma, \tau \rangle$)

A.2.2 Semantic parsetrees

(A.19) A *semantic parsetree* is a structure $\mathcal{T} = (T, \mathbf{D}, V)$, where

a. $T = \{u_1, u_2, \ldots\}$ is a set of nodes.

b. $\mathbf{D}$ is an irreflexive, intransitive binary relation on T; it is defined so that, for all $u \in T$, there is at most one u' such that $\mathbf{D}(u', u)$ and at most two distinct nodes u', u'' such that $\mathbf{D}(u, u')$ and $\mathbf{D}(u, u'')$.

c. $\mathbf{D}^*$, the reflexive, transitive closure of $\mathbf{D}$, is acyclic.

d. There is a unique $r \in T$, the root: there is no $u \in T$ such that $\mathbf{D}(u, r)$.

e. V is a valuation function, taking formulae of $\mathcal{L}_{\text{CI}}$ to sets of nodes in T, according to tree-admissibility conditions (A.14), (A.15), (A.16), (A.17), and (A.18).

Thus, the structures are connected, rooted, acyclic graphs. The branching factor for each node is at most 2, and each node has at most one mother.

A.2.3 Semantic parsetree interpretation

(A.20) **parsetree interpretation**

Let $\mathcal{T}$ be a semantic parsetree with the at-issue term $\alpha : \sigma^a$ on its root node, and distinct terms $\beta_1 : \langle s^a, t^c\rangle,\ \ldots,\ \beta_n : \langle s^a, t^c\rangle$ on nodes in it (extensionally, $\beta_1 : t^c,\ \ldots,\ \beta_n : t^c$). Then the interpretation of $\mathcal{T}$ is the tuple

$$\left\langle [\![\alpha : \sigma^a]\!]^{\mathcal{M}_i,g}, [\![\beta_1 : \langle s^a, t^c\rangle]\!]^{\mathcal{M}_i,g}, \ldots, [\![\beta_n : \langle s^a, t^c\rangle]\!]^{\mathcal{M}_i,g} \right\rangle$$

where $[\![\cdot]\!]^{\mathcal{M}_i,g}$ is the interpretation function, taking formulae of the meaning language to the interpreted structure $\mathcal{M}_i$, relative to a variable assignment g.

A.2.4 Intensional models for $\mathcal{L}_{\mathrm{CI}}$

(A.21) $\mathfrak{M} = \{\mathcal{M}_1, \mathcal{M}_2, \ldots\}$ is a set of intensional models for the logic $\mathcal{L}_{\mathrm{CI}}$. Each $\mathcal{M}_i \in \mathfrak{M}$ is a pair $(\mathbf{D}, V_i)$, where

a. $\mathbf{D}$ is a set of domains, common to all models in $\mathfrak{M}$ and defined as follows:

  i. The domain of e^a and e^c is D_e, a set of entities.

  ii. The domain of s^a and s^c is D_s, a set of entities called worlds, disjoint from D_e.

  iii. The domain of t^a and t^c is $D_t = \{0, 1\}$, the set of truth values.

  iv. The domain of a functional type $\langle\sigma, \tau\rangle$ is $\{f \mid f : D_\sigma \mapsto D_\tau\}$.

  v. The domain of a product type $\langle\sigma \times \tau\rangle$ is $D_{\langle\sigma\times\tau\rangle} = D_\sigma \times D_\tau$.

b. V_i is a valuation taking formulae of $\mathcal{L}_{\mathrm{CI}}$ to the model, constrained so that if $\alpha \in ME_\sigma$, then $V_i(\alpha) \in D_\sigma$.

A.2.5 Interpretation for $\mathcal{L}_{\text{CI}}$

The interpretation function for $\mathcal{L}_{\text{CI}}$ is given by $[\![\cdot]\!]^{\mathcal{M}_i,g}$, where $\mathcal{M}_i$ is an intensional model and g is a variable assignment: if x is a variable of type σ, then $g(x) \in D_\sigma$.

(A.22)

a. $[\![\alpha : \sigma]\!]^{\mathcal{M}_i,g} = \begin{cases} V_i(\alpha) & \text{if } \alpha \text{ is a constant of } \mathcal{L}_{\text{CI}} \\ g(\alpha) & \text{if } \alpha \text{ is a variable of } \mathcal{L}_{\text{CI}} \end{cases}$

b. $[\![\alpha(\beta) : \tau]\!]^{\mathcal{M}_i,g} = [\![\alpha : \langle\sigma,\tau\rangle]\!]^{\mathcal{M}_i,g}([\![\beta : \sigma]\!]^{\mathcal{M}_i,g})$

c. $[\![\lambda x.\, \alpha : \langle\sigma,\tau\rangle]\!]^{\mathcal{M}_i,g} =$ the f such that $f(\mathbf{d}) = [\![\alpha : \tau]\!]^{\mathcal{M}_i,g[x:=\mathbf{d}]}$, where $g[x := \mathbf{d}]$ is an assignment that takes x to $\mathbf{d}$ and maps all variables y distinct from x to $g(y)$.

d. $[\![$**the-speaker**$]\!]^{\mathcal{M}_i,g} = i$

e. $[\![[\alpha \bullet \beta]]\!]^{\mathcal{M}_i,g} = \langle [\![\alpha]\!]^{\mathcal{M}_i,g}, [\![\beta]\!]^{\mathcal{M}_i,g} \rangle$

f. $[\![\neg\alpha]\!]^{\mathcal{M}_i,g} = 1$ iff $[\![\alpha]\!]^{\mathcal{M}_i,g} = 0$

g. $[\![\alpha \vee \beta]\!]^{\mathcal{M}_i,g} = 1$ iff $[\![\alpha]\!]^{\mathcal{M}_i,g} = 1$ or $[\![\beta]\!]^{\mathcal{M}_i,g} = 1$

h. $[\![\exists x[\alpha]]\!]^{\mathcal{M}_i,g} = 1$ iff there is a $\mathbf{d} \in D_\sigma$ such that $[\![\alpha]\!]^{\mathcal{M}_i,g[x:=\mathbf{d}]} = 1$.

i. $[\![\mathsf{G}x[\alpha]]\!]^{\mathcal{M}_i,g} = 1$ iff there are sufficiently many $\mathbf{d} \in D_\sigma$ such that $[\![\alpha]\!]^{\mathcal{M}_i,g[x:=\mathbf{d}]} = 1$ to make this appear lawful.

This definition is of course quite partial. It could be fleshed out in many ways.

A.3 The logic $\mathcal{L}_U$

A.3.1 Syntax of $\mathcal{L}_U$

A.3.1.1 Types for $\mathcal{L}_U$

(A.23)

a. e, t, q, and a are basic types for $\mathcal{L}_U$.

b. $u = q \cup a$

c. If σ and τ are types for $\mathcal{L}_U$, then $\langle\sigma,\tau\rangle$ are types for $\mathcal{L}_U$.

A.3.1.2 Terms for $\mathcal{L}_U$

(A.24) a. $\ulcorner$Aaron aches$\urcorner$, $\ulcorner$Barry bays$\urcorner$, ... are constants of type a.

b. $\ulcorner$Does Aaron ache$\urcorner$, $\ulcorner$Has Barry bayed$\urcorner$, ... are constants of type q.

c. $\ulcorner$**aaron**$\urcorner$, $\ulcorner$**barry**$\urcorner$, ... are constants of type e.

d. $\ulcorner$**utter**$\urcorner$ is a constant of type $\langle u, \langle e, t\rangle\rangle$.

e. $\ulcorner$**frankly**$\urcorner$, $\ulcorner$**honestly**$\urcorner$, ... are constants of type $\langle\langle u, \langle e, t\rangle\rangle, \langle d, t\rangle\rangle$ or $\langle\langle u, \langle e, t\rangle\rangle, \langle q, t\rangle\rangle$.

f. If $\ulcorner\alpha\urcorner$ is a term of $\mathcal{L}_U$ of type $\langle\sigma, \tau\rangle$ and $\ulcorner\beta\urcorner$ is a term of $\mathcal{L}_U$ of type σ, then $\ulcorner\alpha\urcorner(\ulcorner\beta\urcorner)$ is a term of type τ.

g. If $\ulcorner\alpha\urcorner$ is a term of $\mathcal{L}_U$ of type τ, and x is a variable of $\mathcal{L}_U$ of type σ, then $\lambda x.\ulcorner\alpha\urcorner$ is a term of $\mathcal{L}_U$ of type $\langle\sigma, \tau\rangle$.

In analyses stated in $\mathcal{L}_U$, I often employ variables (the conventions are defined just below). However, these are for convenience only. I employ no free variables in $\mathcal{L}_U$.

A.3.1.3 Variable conventions for $\mathcal{L}_U$

u	$\{S, S', S'', \ldots\}$	sentences
d	$\{A, A', A'', \ldots\}$	declaratives
q	$\{Q, Q', Q'', \ldots\}$	questions
$\langle u, \langle e, t\rangle\rangle$	$\{\mathcal{U}, \mathcal{U}', \mathcal{U}'', \ldots\}$	functions from sentences to one-place functions on entities

A.3.1.4 Abbreviated lexicon for $\mathcal{L}_U$

(A.25) $\ulcorner$**perf-hon**$\urcorner \overset{def}{=} \lambda S.\ulcorner$**politely**$\urcorner(\ulcorner$**utter**$\urcorner(S)(\ulcorner$**the-speaker**$\urcorner) : \langle u, t\rangle$

(A.26) *speaking* $\rightsquigarrow \ulcorner$**utter**$\urcorner \overset{def}{=} \lambda S \lambda x.\ulcorner$**utter**$\urcorner(S)(x) : \langle u, \langle e, t\rangle\rangle$ (3.145)

(A.27) *honestly speaking* $\rightsquigarrow$

$$\begin{cases} \lambda S.\ulcorner\textbf{honestly}\urcorner(\ulcorner\textbf{utter}\urcorner(S))(\textbf{the-speaker}) : \\ \qquad\qquad\qquad\qquad\qquad\qquad \langle\langle u, \langle e,t\rangle\rangle, \langle d,t\rangle\rangle \\ \lambda Q.\ulcorner\textbf{honestly}\urcorner(\ulcorner\textbf{utter}\urcorner(\ulcorner\textbf{answer}\urcorner(Q)))(\textbf{the-addressee}) : \\ \qquad\qquad\qquad\qquad\qquad\qquad \langle\langle u, \langle e,t\rangle\rangle, \langle q,t\rangle\rangle \end{cases} \tag{3.153}$$

A.3.2 Discourse structures

(A.28) A *discourse structure* is a tuple $\mathcal{D} = (A, \mathbf{D}, D_u, \mathfrak{M}, \hbar, V_{\mathcal{D}})$, where

a. $A = \{a_1, a_2, \ldots\}$ is a set of discourse participants.

b. $D_u = \{\mathcal{S}_1, \mathcal{S}_2, \ldots\}$ is a set of sentences, the domain of u. Each $\mathcal{S}$ is a pair $(\mathcal{T}^s, \mathcal{T}^m)$, in which $\mathcal{T}^s$ is a syntactic structure and $\mathcal{T}^m$ is its associated semantic parsetree (as defined in (2.36)). D_u contains a subset $D_d = \{\mathcal{D}_1, \mathcal{D}_n, \ldots\}$ of declaratives (the domain of d) and a subset $D_q = \{\mathcal{Q}_1, \mathcal{Q}_n, \ldots\}$ of interrogatives (the domain of q). $D_q \cap D_d = \emptyset$.

c. $\mathbf{D}$ is a set of domains, as defined in (2.61); $A \subseteq D_e$.

d. $\mathfrak{M} = \{\mathcal{M}_1, \mathcal{M}_2, \ldots\}$ is a set of intensional models, as defined in (A.21). All $\mathcal{M}_i \in \mathfrak{M}$ have $\mathbf{D}$ as their set of domains.

e. $\hbar$ is a function that takes each $a_i \in A$ to the model $\mathcal{M}_i \in \mathfrak{M}$, where $\mathcal{M}_i$ can be viewed as the world-view of a_i.

f. $V_{\mathcal{D}}$ is a valuation function, taking constants of $\mathcal{L}_U$ to functions formed from objects in $D_e \cup D_u \cup \{0, 1\}$, constrained so that if α is of type σ, then $V_{\mathcal{D}}(\alpha) \in D_\sigma$.

A.4 Interpretation for $\mathcal{L}_{\mathbf{CI}}$ and $\mathcal{L}_U$

The interpretation function for a discourse structure $\mathcal{D}$ is relativized to a speaker s and an addressee, both members of A, the set of discourse participants for $\mathcal{D}$.

(A.29) a. $[\![\varphi]\!]^{\mathcal{D},s,a} = V_{\mathcal{D}}(\varphi)$ if φ is a formula of $\mathcal{L}_U$.

b. $[\![\mathcal{S}]\!]^{\mathcal{D},s,a} =$ the value of $\mathcal{S}$ determined by (A.20) if $\mathcal{S}$ is a parsetree for $\mathcal{L}_{\mathrm{CI}}$.

Bibliography

Aoun, Joseph and Lina Choueiri. 2000. Epithets. *Natural Language and Linguistic Theory* 18(1):1–39.

Aoun, Joseph, Lina Choueiri, and Norbert Hornstein. 2001. Resumption, movement, and derivational economy. *Linguistic Inquiry* 32(3):371–403.

Asudeh, Ash. 2003. *Resumption as Resource Management*. Ph.D. thesis, Stanford University.

Asudeh, Ash and Richard Crouch. 2002. Derivational parallelism and ellipsis parallelism. In Line Mikkelsen and Christopher Potts, eds., *Proceedings of WCCFL 21*, 1–14. Somerville, MA: Cascadilla Press.

Bach, Kent. 1999. The myth of conventional implicature. *Linguistics and Philosophy* 22(4):367–421.

Barendregt, Henk P. 1992. Lambda calculi with types. In Sampson Abramsky, Dov M. Gabbay, and T. S. E. Maibaum, eds., *Handbook of Logic in Computer Science*, Volume II, 117–309. New York: Oxford University Press. Available at <http://www.cs.kun.nl/~henk/papers.html>.

Barker, Chris and Gina Taranto. 2003. The paradox of asserting clarity. In *WECOL Proceedings*. URL `Available at semanticsarchive.net`.

Barwise, Jon and Robin Cooper. 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy* 4(4):159–219.

Barwise, Jon and John Perry. 1983. *Situations and Attitudes*. Cambridge, MA: MIT Press.

Beaver, David Ian. 1997. Presupposition. In Johan van Benthem and Alice ter Meulen, eds., *Handbook of Logic and Language*, 939–1008. Cambridge, MA and Amsterdam: MIT Press and North-Holland.

Beaver, David Ian. 2001. *Presupposition and Assertion in Dynamic Semantics*. Stanford, CA: CSLI.

Bellert, Irena. 1977. On semantic and distributional properties of sentential adverbs. *Linguistic Inquiry* 8(2):337–351.

Blackburn, Patrick, Maarten de Rijke, and Yde Venema. 2001. *Modal Logic*. Cambridge Tracts in Computer Science 58. Cambridge, UK: Cambridge University Press.

Blackburn, Patrick, Claire Gardent, and Wilfried Meyer-Viol. 1993. Talking about trees. In *Proceedings of the Sixth Conference of the European Chapter of the Association for Computational Linguistics*, 21–29. San Francisco: Morgan Kaufmann.

Blackburn, Patrick and Wilfried Meyer-Viol. 1997. Modal logic and model-theoretic syntax. In Maarten de Rijke, ed., *Advances in Intensional Logic*, 29–60. Dordrecht: Kluwer.

Boeckx, Cedric and Fumikazu Niinuma. 2003. Conditions on agreement in Japanese. *Natural Language and Linguistic Theory* 26.

Boër, Steven E. and William G. Lycan. 1976. *The Myth of Semantic Presupposition*. Bloomington, IN: Indiana University Linguistics Club.

Brody, Michael. 1995. *Lexico-Logical Form: A Radical Minimalist Theory*. Cambridge, MA: MIT Press.

Büring, Daniel. 1999. Topic. In Peter Bosch and Rob van der Sandt, eds., *Focus — Linguistic, Cognitive, and Computational Perspectives*, 142–165. Cambridge, UK: Cambridge University Press.

Büring, Daniel and Katharina Hartmann. 2001. The syntax and semantics of focus-sensitive particles in German. *Natural Language and Linguistic Theory* 19(2):229–281.

Chierchia, Gennaro. 1984. *Topics in the Syntax and Semantics of Infinitives and Gerunds*. Ph.D. thesis, University of Massachusetts, Amherst. [Distributed by GLSA].

Chierchia, Gennaro. 1998. Reference to kinds across languages. *Natural Language Semantics* 6(4):339–405.

Chierchia, Gennaro. 2001. Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface, Ms., University of Milan.

Chierchia, Gennaro and Sally McConnell-Ginet. 1990. *Meaning and Grammar*. Cambridge, MA: MIT Press. Page references to the 2 ed., 2000, MIT Press.

Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.

Cinque, Guglielmo. 1999. *Adverbs and Functional Heads — A Crosslinguistic Perspective*. Oxford: Oxford University Press.

Clements, George. 1975. The logophoric pronoun in Ewe: Its role in discourse. *Journal of West African Linguistics* 10:141–177.

Cooper, Robin. 1983. *Quantification and Syntactic Theory*. Dordrecht: D. Reidel.

Cresswell, M. J. 2002. Static semantics for dynamic discourse. *Linguistics and Philosophy* 25(5–6):545–571.

Cruse, D. A. 1986. *Lexical Semantics*. Cambridge, UK: Cambridge University Press.

Culicover, Peter W. 1992. English tag questions and Universal Grammar. *Lingua* 88(1):21–54.

Dalrymple, Mary, John Lamping, Fernando Pereira, and Vijay Saraswat. 2001. Overview and introduction. In Mary Dalrymple, ed., *Semantics and Syntax in Lexical Functional Grammar*, 1–37. Cambridge, MA: MIT Press.

Dekker, Paul. 2002. A proper architecture for presupposition and quantification, Ms., University of Amsterdam.

Durrell, Martin. 1995. *Hammer's German Grammar and Usage*. Lincolnwood, IL: NTC Publishing Company, 2 ed.

Elbourne, Paul. 2001. E-type anaphora as NP-deletion. *Natural Language Semantics* 9(3):241–288.

Emonds, Joseph. 1976. *A Transformational Approach to English Syntax*. New York: Academic Press.

Emonds, Joseph E. 1979. Appositive relatives have no properties. *Linguistic Inquiry* 10(2):211–243.

Farkas, Donka F. 1997. Evaluation indices and scope. In Anna Szabolsci, ed., *Ways of Scope Taking*, 183–215. Dordrecht: Kluwer.

von Fintel, Kai. 1993. Exceptive constructions. *Natural Language Semantics* 1(2):123–148.

von Fintel, Kai. 1994. *Restrictions on Quantifier Domains*. Ph.D. thesis, University of Massachusetts, Amherst, MA. Reproduced and distributed by the GLSA, UMass, Amherst, Linguistics Department.

von Fintel, Kai. 2001. Would you believe it? the king of france is back! (presuppositions and truth-value intuitions). In Anne Bezuidenhout and Marga Reimer, eds., *Descriptions and Beyond: An Interdisciplinary Collection of Essays on Definite and Indefinite Descriptions and other Related Phenomena*. Oxford: Oxford University Press.

Fox, Danny. 2000. *Economy and Semantic Interpretation*. Cambridge, MA: MIT Press.

Francescotti, Robert M. 1995. *Even*: the conventional implicature approach reconsidered. *Linguistics and Philosophy* 18(1):153–173.

Frazier, Lyn. 2003. Processing silence, workshop on Ellipsis, UC Santa Cruz.

Gamut, L. T. F. 1991. *Logic, Language, and Meaning*, Volume 1. Chicago: University of Chicago Press.

Gawron, Jean Mark. 1996. Quantification, quantificational domains, and dynamic logic. In Shalom Lappin, ed., *The Handbook of Contemporary Semantic Theory*, 247–267. Oxford: Blackwell Publishers.

Gazdar, Gerald. 1979a. *Pragmatics: Implicature, Presupposition and Logical Form*. New York: Academic Press.

Gazdar, Gerald. 1979b. A solution to the projection problem. In Oh and Dinneen (1979), 57–89.

Geurts, Bart. 1998. The mechanisms of denial. *Language* 74(2):274–307.

Green, Mitchell S. 2000. Illocutionary force and semantic content. *Linguistics and Philosophy* 23(5):435–473.

Grice, H. Paul. 1975. Logic and conversation. In Peter Cole and Jerry Morgan, eds., *Syntax and Semantics*, Volume 3: Speech Acts, 43–58. New York: Academic Press.

Groenendijk, Jeroen. 1999. The logic of interrogation. In Tanya Matthews and Devon Strolovitch, eds., *Proceedings of SALT IX*, 109–126. Ithaca, NY: Cornell University.

Gunlogson, Christine. 2001. *True to Form: Rising and Falling Declaratives as Questions in English*. Ph.D. thesis, UC Santa Cruz, Santa Cruz, CA.

Halvorsen, Per-Kristian and William A. Ladusaw. 1979. Montague's 'Universal grammar': An introduction for the linguist. *Linguistics and Philosophy* 3(2):185–223.

Harada, S. I. 1976. Honorifics. In Masayoshi Shibatani, ed., *Syntax and Semantics*, Volume 5: Japanese Generative Grammar, 499–561. New York: Academic Press.

Hawkins, John A. 1991. On (in)definite articles: Implicatures and (un)grammaticality prediction. *Journal of Linguistics* 27:405–442.

Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachussets, Amherst. [Distributed by GLSA].

Heim, Irene. 1983. On the projection problem for presuppositions. In Michael Barlow, Daniel P. Flickinger, and Michael T. Wescoat, eds., *Proceedings of WCCFL 2*, 114–125. Stanford, CA: Stanford Linguistics Association.

Heim, Irene. 1992. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics* 9(2):183–221.

Heim, Irene and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell Publishers.

Herzberger, Hans G. 1973. Dimensions of truth. *Journal of Philosophical Logic* 2(4):535–556.

Heycock, Caroline and Anthony Kroch. 1999. Pseudocleft connectivity: implications for the LF interface level. *Linguistic Inquiry* 30(3):365–397.

Higgins, Roger Francis. 1973. *The Pseudo-cleft Construction in English*. Ph.D. thesis, MIT. [Published by Garland, 1979].

Hoeksema, Jacob. 1995. The semantics of exception phrases. In Jaap van der Does and Jan van Eijck, eds., *Quantifers, Logic, and Language*, 145–177.

Horn, Laurence R. 1989. *A Natural History of Negation*. Chicago: University of Chicago Press. Reissued 2001 by CSLI.

Horn, Laurence R. 1991. When new content isn't. *Journal of Pragmatics* 15(4):313–336.

Huddleston, Rodney and Geoffrey K. Pullum. 2002. *The Cambridge Grammar of the English Language*. Cambridge, UK: Cambridge University Press.

Jackendoff, Ray S. 1972. *Semantic Interpretation in Generative Grammar*. Cambridge, MA: MIT Press.

Jackendoff, Ray S. 1984. On the phrase "the phrase *the phrase*". *Natural Language and Linguistic Theory* 2(1):25–37.

Jacobson, Pauline. 1999. Towards a variable-free semantics. *Linguistics and Philosophy* 22(2):117–184.

Jacobson, Pauline. 2000. Paycheck pronouns, Bach-Peters sentences, and variable-free semantics. *Natural Language Semantics* 8(2):77–155.

Johnson, David E. and Paul M. Postal. 1980. *Arc Pair Grammar*. Princeton, NJ: Princeton University Press.

Kamp, Hans. 1981. A theory of truth and discourse representation. In Jeroen Groenendijk, Theo M. V. Janssen, and Martin Stockhof, eds., *Formal Methods in the Study of Language*, 277–322. Amsterdam: Mathematical Centre.

Kamp, Hans and Uwe Reyle. 1993. *From Discourse to Logic. Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Dordrecht: Kluwer.

Kaplan, David. 1989. Demonstratives: an essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In Joseph Almog, John Perry, and Howard Wettstein, eds., *Themes from Kaplan*, 481–614. New York: Oxford University Press. [Versions of this paper began circulating in 1971].

Kaplan, David. 1999. What is meaning? Explorations in the theory of *Meaning as Use*, brief version — Draft 1.

Karttunen, Lauri. 1971. Implicative verbs. *Language* 47(2):340–358.

Karttunen, Lauri. 1973. Presuppositions and compound sentences. *Linguistic Inquiry* 4(2):169–193.

Karttunen, Lauri. 1976. Discourse referents. In James D. McCawley, ed., *Syntax and Semantics*, Volume 7: Notes from the Linguistic Underground, 363–385. New York: Academic Press.

Karttunen, Lauri and Stanley Peters. 1979. Conventional implicature. In Oh and Dinneen (1979), 1–56.

Katz, Fred M. and Jeffrey J. Katz. 1977. Is necessity the mother of intension? *The Philosophical Review* 86(1):70–96.

Keenan, Edward L. 1971. Two kinds of presupposition in natural language. In Charles J. Fillmore and D. Terence Langendoen, eds., *Studies in Linguistic Semantics*, 45–54. New York: Holt, Rinehart, Winston.

Kempson, Ruth. 2003a. Nonrestrictive relatives, and the growth of logical form, Ms., King's College London.

Kempson, Ruth. 2003b. Nonrestrictive relatives, and the growth of logical form. In *Proceedings of WCCFL 22*. Somerville, MA: Cascadilla Press.

Kennedy, Christopher and Jason Merchant. 2000. Attributive comparative deletion. *Natural Language and Linguistic Theory* 18(1):89–146.

Klein, Ewan and Ivan A. Sag. 1985. Type-driven translation. *Linguistics and Philosophy* 8(2):163–201.

Krahmer, Emiel. 1998. *Presupposition and Anaphora*. Stanford, CA: CSLI.

Kratzer, Angelika. 1999. Beyond *ouch* and *oops*: How descriptive and expressive meaning interact. A comment on David Kaplan's paper, Cornell Conference on Theories of Context Dependency (March 26th). <http://semanticsarchive.net/Archive/WEwNGUyO/>.

Lakoff, George. 1966. Deep and surface grammar, Ms., Harvard University.

Lakoff, George. 1970. Global rules. *Language* 46:627–639.

Lasnik, Howard. 1976. Remarks on coreference. *Linguistic Analysis* 2(1):1–22. Reprinted in Lasnik (1989), 90–109.

Lasnik, Howard. 1989. *Essays on Binding and Anaphora*. Dordrecht: Kluwer.

Löbner, Sebastian. 2002. *Understanding Semantics*. London: Arnold.

Loetscher, Andreas. 1973. On the role of nonrestrictive relative clauses in dicourse. In Claudia Coruma, T. Cedric Smith-Stark, and Ann Wesler, eds., *Papers from the Ninth Regional Meeting of the Chicago Linguistic Society*, 356–368. Chicago: CLS.

McCawley, James D. 1982. Parentheticals and discontinuous structure. *Linguistic Inquiry* 13(1):91–106.

McCawley, James D. 1987. Some additional evidence for discontinuity. In Geoffrey J. Huck and Almerindo E. Ojeda, eds., *Syntax and Semantics*, Volume 20: Discontinuous Constituency, 185–200. Orlando, FL: Academic Press.

McCawley, James D. 1989. Individuation in and of syntactic structures. In Mark R. Baltin and Anthony S. Kroch, eds., *Alternative Conceptions of Phrase Structure*, 117–138. Chicago: University of Chicago Press.

McCawley, James D. 1998. *The Syntactic Phenomena of English*. Chicago: Chicago University Press, 2 ed.

Merchant, Jason. 2002. PF output constraints and elliptical repair in SAI comparatives. In Line Mikkelsen and Christoper Potts, eds., *Proceedings of WCCFL 21*, 292–305. Somerville, MA: Cascadilla Press.

Merchant, Jason. 2003. Fragments and ellipsis, Ms., University of Chicago.

Mikkelsen, Line. 2002a. Dissertation prospectus, uC Santa Cruz.

Mikkelsen, Line. 2002b. Specification is not inverted predication. In Mako Hirotani, ed., *Proceedings of NELS 32*, 403–422. Amherst, MA: GSLA.

Mikkelsen, Line. 2002c. Two types of definite description subjects. In Malvina Nissim, ed., *Proceedings of The Student Session at ESSLLI 14*, 141–153. Trento, Italy.

Mikkelsen, Line. 2003. *Specifying Who: The Structure, Meaning, and Use of Specificational Clauses*. Ph.D. thesis, UC Santa Cruz.

Moltmann, Frederike. 1995. Exception sentences and polyadic quantification. *Linguistics and Philosophy* 18(3):223–80.

Montague, Richard. 1970a. English as a formal language. In Bruno Visentini et al., ed., *Linguaggi nella società e nella tecnica*, 189–224. Milan: Edizioni di Communità. Reprinted in Montague (1974), 188–221.

Montague, Richard. 1970b. Universal grammar. *Theoria* 36:373–398. Reprinted in Montague (1974), 222–246.

Montague, Richard. 1973. The proper treatment of quantification in ordinary English. In Jaakko Hintikka, Julius Matthew Emil Moravcisk, and Patrick Suppes, eds., *Approaches to Natural Language*, 221–242. Dordrecht: D. Reidel. Reprinted in Montague (1974), 247–270.

Montague, Richard. 1974. *Formal Philosophy: Selected Papers of Richard Montague*. New Haven, CT: Yale University Press.

Nilsen, Øystein. 2001. Adverb order in type-logical grammar. In Rob van Rooy and Martin Stokhof, eds., *Proceedings of the Amsterdam Colloquium*.

Nilsen, Øystein. 2003. Domains for adverbs. *Lingua* .

Nunberg, Geoffrey. 1990. *The Linguistics of Punctuation*. Stanford, CA: CSLI.

Oh, Choon-Kyu and David A. Dinneen, eds. 1979. *Syntax and Semantics*, Volume 11: Presupposition. New York: Academic Press.

Partee, Barbara. 1987. Noun phrase interpretation and type-shifting principles. In Jeroen Groenendijk, Dick de Jong, and Martin Stokhof, eds., *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers*, 115–143. Dordrecht: Foris Publications. Reprinted in Portner and Partee (2002), 357–381.

Partee, Barbara, Alice ter Meulen, and Robert E. Wall. 1993. *Mathematical Methods in Linguistics*. Corrected 1st edition. Dordrecht: Kluwer.

Partee, Barbara H. 1999. Weak NP's in HAVE sentences. In Maarten Marx. Maarten de Rijke Jelle Gerbrandy and Yde Venema, eds., *JFAK: Essays dedicated to Johan van Bentham on the Occasion of his 50th Birthday*. Vossiuspers: Amsterdam University Press.

Partee, Barbara Hall and Mats Rooth. 1983. Generalized conjunction and type ambiguity. In Rainer Bäuerle, Christoph Schwarze, and Armin von Stechow, eds., *Meaning, Use, and Interpretation of Language*, 361–383. Berlin: Walter de Gruyter.

Percus, Orin and Uli Sauerland. 2002. Binding restrictions in dream reports. Handout. NELS 33, MIT (November 8–10).

Portner, Paul and Barbara H. Partee, eds. 2002. *Formal Semantics: The Essential Readings*. Oxford: Blackwell Publishing.

Postal, Paul M. 1998. *Three Investigations of Extraction*. MIT Press: Cambridge, MA.

Potts, Christopher. 2002a. The lexical semantics of parenthetical-*as* and appositive-*which*. *Syntax* 5(1):55–88.

Potts, Christopher. 2002b. The syntax and semantics of *As*-parentheticals. *Natural Language and Linguistic Theory* 20(3):623–689.

Potts, Christopher. 2003a. Expressive content as conventional implicature. In Makoto Kadowaki and Shigeto Kawahara, eds., *Proceedings of the North East Linguistics Society 33*. Amherst, MA: GLSA.

Potts, Christopher. 2003b. A layered semantics for utterance modifiers, Talk presented at the Workshop on Direct Compositionality, Brown University, June 19–21, 2003.

Potts, Christopher and Geoffrey K. Pullum. 2003. Model theory and the content of OT constraints. *Phonology* 19(1):1–33.

Pullum, Geoffrey K. 1979. *Rule Interaction and the Organization of Grammar*. Outstanding Dissertations in Linguistics. New York: Garland.

Quang, Phuc Dong. 1971. English sentences without overt grammatical subject. In Arnold M. Zwicky, Peter H. Salus, Robert I. Binnick, and Anthony L. Vanek, eds., *Studies out in Left Field: Defamatory Essays Presented to James D. McCawley on the Occasion of his 33rd or 34th birthday*, 3–10. Amsterdam: John Benjamins.

Reynolds, John C. 1983. Types, abstraction and parametric polymorphism. In R. E. A. Mason, ed., *Information Processing 83: Proceedings of the IFIP 9th World Computer Congress*, 513–523. Amsterdam: Elsevier Science.

Rogers, James. 1998. *A Descriptive Approach to Language-Theoretic Complexity*. Stanford, CA: CSLI/FoLLI.

Rooth, Mats. 1985. *Association with Focus*. Ph.D. thesis, University of Massachusetts, Amherst.

Rooth, Mats. 1992. A theory of focus interpretation. *Natural Language Semantics* 1(1):75–116.

Ross, John Robert. 1967. *Constraints on Variables in Syntax*. Ph.D. thesis, MIT, Cambridge, MA. Published as *Infinite Syntax!*. Norwoord, NJ: Ablex (1986).

Ross, John Robert. 1973. Slifting. In Maurice Gross, Morris Halle, and Marcel-Paul Schützenberger, eds., *The Formal Analysis of Natural Languages*, 133–169. The Hague: Mouton de Gruyter.

Sampson, Geoffrey. 1975. The single mother condition. *Journal of Linguistics* 11(1):1–11.

Sauerland, Uli. 2001. On the computation of conversational implicatures. In Rachel Hastings, Brendan Jackson, and Zsofia Zvolenszky, eds., *Proceedings of SALT 11*. Ithaca, NY: Cornell Linguistics Circle.

Sauerland, Uli and Fabian Heck. 2002. LF-intervention effects in pied-piping. Handout. NELS 33, MIT (November 8–10).

Schlenker, Philippe. 2003. A plea for monsters. *Linguistics and Philosophy* 26(1):29–120.

Shan, Chung-chien. 2001. A variable-free dynamic semantics. In Robert van Rooy and Martin Stokhof, eds., *Proceedings of the 13th Amsterdam Colloquium*, 204–209. Amsterdam: Institute for Logic, Language, and Computation.

Shan, Chung-chien. 2002. A continuation semantics of interrogatives that accounts for Baker's ambiguity. In Brendan Jackson, ed., *Proceedings of SALT 12*, 246–265. Ithaca, NY: CLC Publications.

Sharvit, Yael. 1999. Connectivity in specificational sentences. *Natural Language Semantics* 7(3):299–346.

Shibatani, Masayoshi. 1978. Mikami Akira and the notion of 'subject' in Japanese grammar. In J. Hinds and I. Howards, eds., *Problems in Japanese Syntax and Semantics*. Kaitakusha.

Soames, Scott. 1989. Presupposition. In Dov M. Gabbay and Franz Guenthner, eds., *Handbook of Philosophical Logic*, Volume IV: Topics in the Philosophy of Language, 553–616. Dordrecht: Kluwer.

Soames, Scott. 2002. *Beyond Rigidity*. Oxford: Oxford University Press.

Stalnaker, Robert. 1979. Assertion. In Peter Cole, ed., *Syntax and Semantics*, Volume 9: Pragmatics, 315–332. New York: Academic Press. Reprinted in Stalnaker (1999), 78–95.

Stalnaker, Robert C. 1999. *Context and Content*. Oxford Cognitive Science. Oxford: Oxford University Press.

von Stechow, Arnim. 2002. Binding by verbs: Tense, person and mood under attitudes, Presented at NELS 33, MIT. Available at <http://vivaldi.sfs.nphil.uni-tuebingen.de/~arnim10/>.

Steedman, Mark. 2000. Information structure and the syntax–phonology interface. *Linguistic Inquiry* 31(4):649–689.

Stump, Gregory. 1981. The interpretation of frequency adjectives. *Linguistics and Philosophy* 5(2):221–256.

Stump, Gregory T. 1985. *The Semantic Variability of Absolute Constructions*. Dordrecht: D. Reidel.

Thorne, James Peter. 1972. On nonrestrictive relative clauses. *Linguistic Inquiry* 3(4):552–556.

Toribio, Almeida Jacqueline. 1990. Specifier-head agreement in Japanese. In Aaron L. Halpern, ed., *The Proceedings of the Ninth West Coast Conference on Formal Linguistics*, 535–548. Stanford Linguistics Association, Stanford, CA: CSLI.

van Benthem, Johan. 1991. *Language in Action: Categories, Lambdas, and Dynamic Logic*. Amsterdam: North-Holland.

van der Sandt, Rob. 1988. *Context and Presupposition*. London: Croom Helm.

van der Sandt, Rob. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* 9(4):333–377.

Williams, Joseph M. 1990. *Style: Toward Clarity and Grace*. Chicago: University of Chicago Press. With two chapters coauthored by Gregory G. Colomb.

Winter, Yoad. 2002. *Flexibility Principles in Boolean Semantics*. Cambridge, MA: MIT Press.

Zimmermann, Malte. 2000. Pluractional quantifiers: the *occasional*-construction in English and German. In Brendan Jackson and Tanya Matthews, eds., *Proceedings of SALT X*, 290–306. Ithaca, NY: Cornell University.